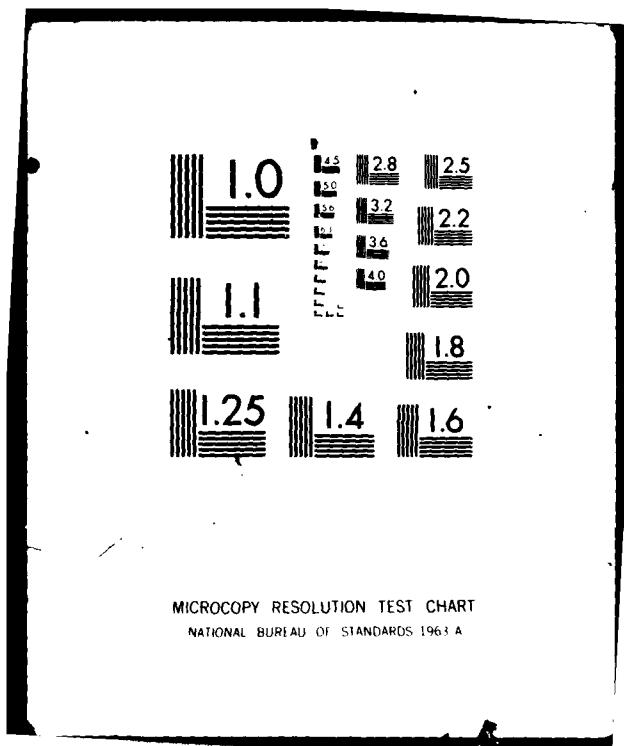


AD-A097 267 ADMIRALTY MARINE TECHNOLOGY ESTABLISHMENT TEDDINGTON--ETC F/G 5/8
APU MAN-COMPUTER STUDIES GROUP SEMINAR PROCEEDINGS I FEBRUARY 1--ETC(U
JUN 80 R GREGORY
UNCLASSIFIED AMTE(E)TM80190 DRIC-BR-76709 NL

END
DATE
1980
3-8h
DTIC



APU MAN-COMPUTER STUDIES GROUP

SEMINAR PROCEEDINGS I
FEBRUARY 1980 - APRIL 1980

BY

R GREGORY

Accession For	
NTIS	GR&I
DTIC TAB	
Unannounced	
Justification	
By	
Distribution	
Availability Codes	
Dist	Avail and/or Special

A

Summary

The Man Computer Studies Group was set up at AMTE/APU in January 1980 as an informal working group to explore the potential of computers as intelligent aiding devices in both training and operational contexts.

The three papers that constitute this document have crystallised from a series of Group seminars held between February and April 1980.

Paper I offers a cybernetic framework for understanding computer based learning and focuses on the importance of subject matter or knowledge representation. Papers II and III discuss two techniques of subject matter representation emanating from the laboratories of Ira Goldstein of MIT and Gordon Pask of SRL, Richmond respectively.

The three papers, together, constitute volume one of the Group's seminar proceedings.

AMTE (Teddington)
Queen's Road
TEDDINGTON Middlesex TW11 0LN

June 1980

(C)

Copyright
Controller HMSO London
1980

52pp
19 figs

CONTENTS

INTRODUCTION

PAPER I. THREE COMPUTER ROLES IN A CYBERNETIC LEARNING SYSTEM

1. Introduction
2. Cybernetic Learning Systems
 - 2.1 The Learner
 - 2.2 Subject Matter Representation
 - 2.3 The Subject Matter Expert
 - 2.4 The Tutor
3. Three Different Computer Roles in Computer Aided Learning
 - 3.1 Pre-CAL
 - 3.2 CAL
 - 3.3 Intelligent CAL
4. Conclusion
- References

PAPER II. THE REPRESENTATION OF KNOWLEDGE: GOLDSTEIN'S GENETIC GRAPH

1. Introduction
2. Expert-Based vs Learner-Based Modelling
3. The Genetic Graph
4. Problems
5. Conclusions
- References
- Appendix

PAPER III. THE REPRESENTATION OF KNOWLEDGE: PASK'S Lp

1. Introduction
2. Lp: Operations and Operational concepts
 - 2.1 Derivation, Coherency and Organisational Closure
 - 2.2 The Rule of Genoa and Analogy
 - 2.3 Pruning, Selective Pruning and Superimposition
 - 2.4 Saturation
 - 2.5 Condense and Analogy
 - 2.6 Isomorphic Inference
3. Summary
- References

GENERAL REFERENCES

GENERAL INTRODUCTION

The Man Computer Studies Group was set up at APU in January 1980 as an informal working group to explore the potential of computers as intelligent aiding devices in both training and operational contexts.

During the period February-April 1980, the group participated in several seminars designed to cover important areas of interest. The three papers that constitute this document are crystallisations of those seminars and are presented as Volume One of the Group's Seminar Proceedings.

The main prerequisites to the seminars themselves involved the identification and collection of those references considered most relevant to the enquiries and a flavour of APU's interests and emphases may be ascertained by consulting the General Reference section at the end of this document. The task of establishing the references was undertaken by Sheppard and Gregory and copies were disseminated to the other participants under the rubric 'Essential Reading'. One of the main purposes of these seminars was to develop familiarity with the material and to construct a shared understanding of the topics discussed. It is in this context that the formalisation of these seminars - the papers - should be considered.

The first of these papers offers a cybernetic framework for understanding computer based learning. Four components of such a learning system are identified and then three different distributions of these components are discussed. These three distributions describe three different computer roles in learning systems.

The second and third papers are concerned with one of the components discussed in the first paper - that of subject matter representation - and the subject matter, or knowledge representation schemes of Ira Goldstein of MIT and Gordon Pask of SRL, Richmond are explored in successive papers. In our view, the problem of knowledge representation is crucial in the attempt to construct intelligent computer systems. Here, 'intelligent' means that computers are enabled to adapt themselves to the styles and requirements of their users: that they are made sensitive to operating characteristics of the particular individuals using them.

One of the reasons why people can act intelligently is that they have sentience. That is, people do not just have knowledge, but they know what it is to have knowledge. A major problem that has faced researchers in Artificial Intelligence (AI) is that machines are quite simply not sentient, and thus to coerce machines into appearing to adapt to ourselves as users has required the simulation of sentience. The field of AI requires computers not just to remember what knowledge they have - this is comparatively trivial - but to appear to understand what it means to have particular knowledge at particular times.

We do not know how we are sentient but attempts have been made to

*Participants: C Sheppard, R Gregory, M Rowley, E Wheatley, E Porteous, R Todd, D Cunningham.

simulate sentience in computers by representing knowledge in terms of some connectivity. Different systems of connectivity have included the epistemological links of Goldstein's Genetic Graph, Pask's coherency networks and the chains, hierarchies and heterarchies of other researchers. The point here is that computers cannot understand anything, but understanding can be simulated, to some degree at least, by devising some system of knowledge connectivity that is rich enough to allow the machine to suggest significant but non-obvious associations that have meaning for the user. It is evident that the success of a knowledge representation system, and hence of intelligent computer aided training and operating, is to a large extent dependent on the particular system of connectivity that the technique embodies and two such connectivities are the subject of Papers II and III.

PAPER I: THREE COMPUTER ROLES IN A CYBERNETIC LEARNING SYSTEM

INTRODUCTION

One of the more useful ways of understanding the computer's role in training is to conceptualise the learning process as cybernetic: that is to say, learning is controlled. This control may exist at several system loci and the amount of control at each locus may vary throughout the course of the learning session. These loci are the Learner himself, the Representation(s) of the Expert(s) (Expert Models) and the teaching aids incorporated into the system. All of these system components may be viewed as competing to control the interaction between Learner and task. The control is mediated by a Tutor characterised by one or more teaching strategies together with their management, and the location of the system Tutor represents one of the more important factors in deciding what the role of the computer actually is.

2. CYBERNETIC LEARNING SYSTEMS

To understand what a cybernetic learning system is requires four main notions and these are: the Learner; Subject Matter Representation; Subject Matter Expert; and Tutor.

(a) The Learner.

The main point to be made here is that the Learner may or may not be a sophisticated Learner (see Augstein & Thomas, 1976) that is to say, he may be more or less aware of what makes learning efficient for him. The extent to which he has learnt how to learn will have implications for the amount and type of explicit instruction he will require - largely irrespective of subject matter. For example, the Learner may have evolved efficient strategies for scanning text and picking out the relevant points or he may not, in which case his attention must be externally directed, if learning is to be made efficient.

(b) Subject Matter Representation.

It is important to understand the technical meaning of this term. A representation has existence and it is brought into existence in order to fulfil a particular purpose or purposes. It organises a domain. A task domain consists of every conceivable and hypothetical aspect of a task - a task's entirety. A domain is an amorphous whole and does not exist in any concrete form until it is represented in a particular way for particular purposes. Many subject matter representations may be derived from one task domain and they may differ along several dimensions, including their grain of detail (the fineness of distinction between concepts or nodes in the representation), the richness of their connective structure (what the links between components or concepts allow), their medium of existence (hard/software, pencil and paper, actual performance, verbal description etc), the number of purposes that they have the potential to achieve and the nature of these purposes.

Bainbridge's (1979) concern that some verbal data are misleading when compared with "observed behaviour" whilst others are not, is

unnecessary. Different representations (ie organisations) of the domain satisfy different purposes and these should be made explicit when the method of representation is selected.

Some concrete examples of subject matter representations within the task domain of Fighter Controlling (see Court & Brooking 1978, and Gregory 1979b) will serve to clarify things.

- (i) A high fidelity simulation of a Fighter Controller's radar is a subject matter representation expressed in the medium of computer hardware and software. This representation is fine grained but has an impoverished connective structure because its components must be linked together in the particular way the designer intended. Failure to comply with this design results in 'malfunction'. In spite of this representation's connective structure, it may nevertheless serve several different purposes: although initially expensive, it is cost effective compared with the equivalent hours spent on live control; it is safer than live control; etc.
- (ii) A schematic diagram of a Fighter Control Simulator is a representation expressed in the medium of paper and pencil. This representation is coarse grained and it too has an impoverished connective structure for the same reason as at (i). Its purpose might be to provide an overall view of the system to interested parties.
- (iii) A computer implementation of an Expert Model is also a representation of the task domain. It is fine grained (or must be if it is to be actionable) and its medium of existence is software. The richness of its connective structure will depend, however, on the technique used to derive the representation: that is, on how far the representational technique embodies a sophisticated connectivity. An Expert Model based on behavioural chains and contingencies (eg Behavioural Objectives Analysis), simple hierarchies of action or process (eg Gagne 1960) or behavioural taxonomies (eg Folley 1962a, Miller 1962, etc) have relatively impoverished connectivities compared with the Genetic Graph and l_p techniques of Goldstein (1979) and Pask (1979, and in press) respectively.* A major consequence of impoverished connectivity in this type of representation is that individual differences in performance cannot be modelled (represented). Clearly if the technique only allows one way of describing a task or sub task (or only one way of making it actionable) then it fails to represent the different but equally valid ways produced by different Experts, and by the same Expert on different occasions. Perhaps the main point is that such a technique fails to hold up as important that individuals do have preferred ways of thinking about a task and differ from each other in this respect.

*See this volume for discussion of both techniques.

(c) The Subject Matter Expert

The important, though obvious, point about the Subject Matter Expert is that he "knows" and structures his task and his relationship with the task in ways that allow him to produce expert performance (see Gregory 1979a). Now, we can seek to represent this internal structure by constructing an Expert Model. But we should recall that such an Expert Model is a particular sort of subject matter representation produced for particular purposes (schematically, to give ourselves an idea of task entailments, or on a computer to simulate expert performance etc). We should not make the mistake, therefore, of considering an Expert Model and the Subject Matter Expert from whom it was elicited as the same thing: one merely represents the other. Any Expert Model we do elicit, by whatever technique, is done so for our purposes and we cannot conclude that the representation is experienced by the Expert. It may seem, in consequence, that the Subject Matter Expert may be represented in different ways for different purposes, thus acquiring for himself, domain status. This is not quite right, however. Actual performance, questionnaire replies, interview data, and debriefings generated by the same Subject Matter Expert are all different representations of one particular perspective (ie. the Subject Matter Expert) on the task domain. The notion of perspective is taken up again in the paper on Lp.

(d) The Tutor

The function of the Tutor is, as we have already said, to mediate learning. An efficient Tutor is sensitive to the Learner's style of learning and to the extent of the Learner's knowledge. Style of learning is matched with the Tutor's teaching (which, naturally, exists as a representation of the teaching domain), while the Learner's knowledge state is compared with an Expert Model. Hence the Tutor is able to decide what to teach at any one time and how to teach it. In fact, the first decision the Tutor must make is whether to intervene at all. The outcome of this decision depends on the perspective on the teaching domain that is represented in the Tutor. The perspective that we are taking at APU is characterised by supportive considerations rather than concern to indoctrinate, (Sheppard, in press, Sheppard, Gregory, Rowley in press).

In the cybernetic view of things then, we start with a Learner interacting with some computer simulated task environment (subject matter representation). This interaction is analysed on-line to construct a model of the Learner in terms of both knowledge state and learning style (subject matter representations from the perspective on the task domain afforded by the Learner). These Learner Models are compared with one or more Expert Models (subject matter representations from the perspective(s) on the task domain afforded by the Subject Matter Expert(s)) and, following a decision to intervene, the Tutor seeks to encourage the growth of the Learner's knowledge state via the selected teaching strategy and appropriate teaching aids.

3. THREE DIFFERENT COMPUTER ROLES

Thus far we have discussed the notions involved in understanding what a cybernetic learning system is. In the rest of this paper, three major roles that a computer may play in a learning system are discussed from a cybernetic point of view. In particular, the location of the system Tutor - the component that mediates the learning process - will be shown to be directly related to the apparent intelligence of the system. Here intelligence is used to mean the sensitivity of the computer to the particular individual who is designated Learner.

(a) Pre-Computer Aided Learning (Pre-CAL)

Fig 1 illustrates what may be termed a Pre-CAL system and it is this system that is traditionally employed by the RN.

Typically, the equipment represents the task domain by providing a high fidelity simulation of the operator's station. The main point to be made about this system is that the major part of it is controlled by the (human) Instructor. Thus the Instructor is both Subject Matter Expert and Tutor. As Tutor, he may construct representations of the Learner by attending to what the Learner says and does, although the nature of the learning model (as distinct from knowledge model) will depend, in part, on the way the Instructor is predisposed to believe that people learn. Ideally, these models (usually unarticulated) are used for two purposes: first to compare with his own knowledge of the task as a Subject Matter Expert, and second, to match the way in which the student is acquiring this knowledge with an appropriate tutorial method and teaching aids. In practice, however, we may expect these processes to be less than optimal - especially the second, since the Navy Instructor achieves that status by virtue of his being a Subject Matter Expert, rather than because he has access to rich personal tutorial resources.

From Fig 1 it is clear that the role of the computer is nothing more than to simulate the task environment. Its instructional role is severely restricted to providing the means for a Subject Matter Expert to demonstrate particular points or exercises that the Tutor deems worthwhile.

(b) CAL

Fig 2a illustrates a CAL System.

It is a CAL system because the Tutor has access to a repertoire of teaching aids that includes computer-based devices designed to improve the quality of feedback to the Learner. The computer is not an intelligent system component however, because the decision to select these aids, according to some teaching strategy, on particular occasions is made by a Tutor who is human. Computer based aids include, in the case of the Fighter Control Skills Trainer currently under construction at APU (Gregory, in press), simple explanations to the student of, eg why the simulation has stopped, air picture enhancements and emphases, record and replay facilities and computer driven illustrations of intercept solutions, the latter derived from a 'black box' or mathematical model of the intercept problem.

 = INSTRUCTOR
 = COMPUTER
 = STUDENT

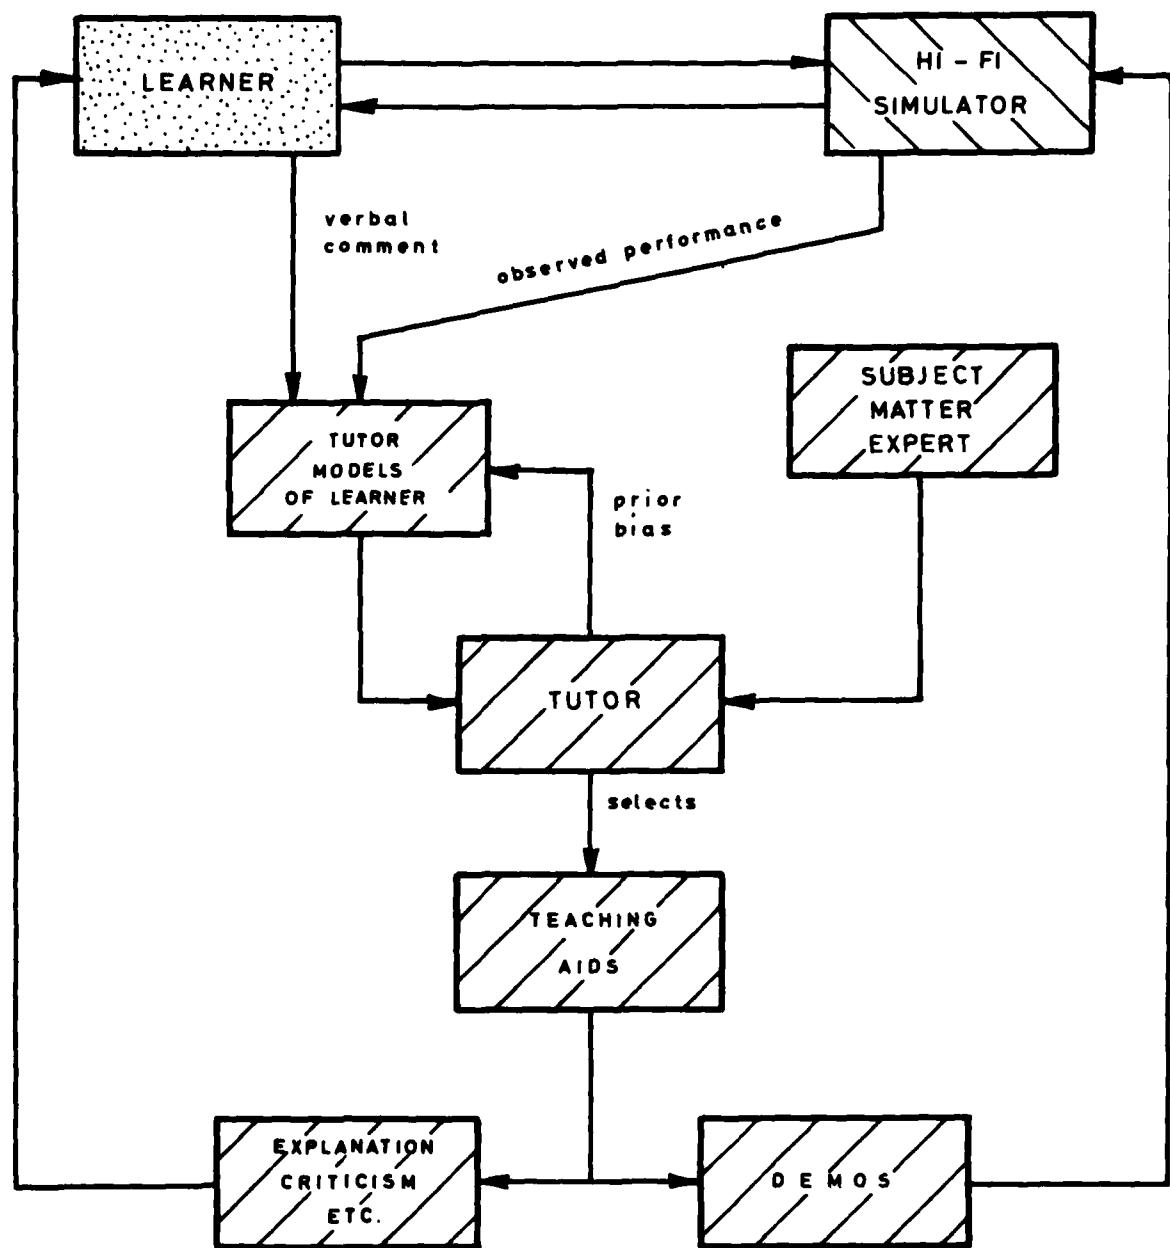


FIG. 1 PRE - CAL : THE TRADITIONAL RN USE OF COMPUTERS FOR TRAINING

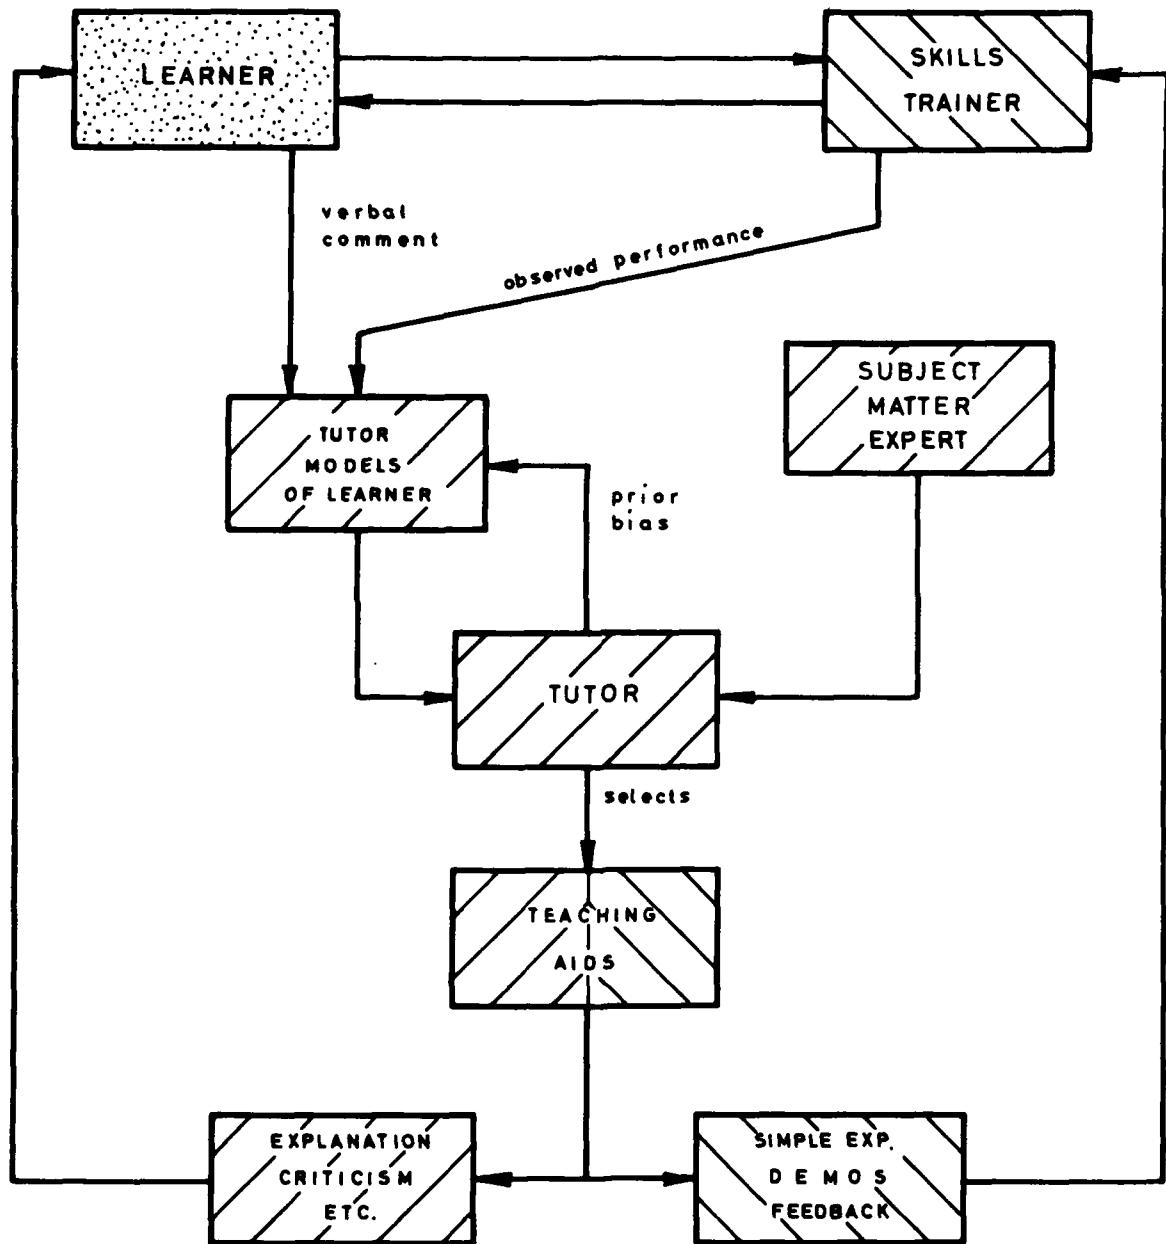
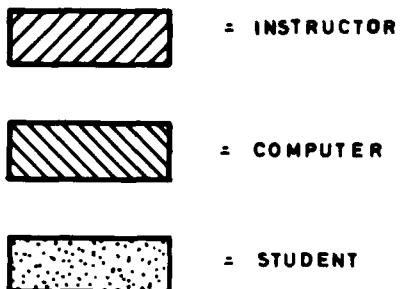


FIG. 2a CAL : STUDENT AND INSTRUCTOR

(See Burton & Brown, 1979 for discussion of 'black box' and 'glass box' experts).

Figure 2b illustrates a CAL variation that pertains when the Instructor is removed from the system. This alternative student-only-mode will be available on the aforementioned Fighter Control Skills Trainer. In this case, the same computer based aids are available but the Learner decides when and how to use them. The Learner is therefore teaching himself, and comparing the state and structure of his own knowledge with what he imagines the state and structure of the Subject Matter Expert's knowledge to be. Actually just as one source of error is reduced, another is increased. Since the Tutor is now "inside" the Learner, there is now little possibility of a mismatch between the Learner's actual understanding of the task and the Tutor's model of that understanding. (The student's learning efficiency will depend, of course, on how far he is aware of his own learning process). However, error is likely to be introduced which is inversely proportional to the correspondence between the Learner's imagined Subject Matter Expert, and the Subject Matter Expert himself. It is easy to conceive of a situation where the Learner efficiently teaches himself to execute a particular intercept type which the Subject Matter Expert would never execute because it conflicts with some higher order constraint that the Learner was simply not aware of.

(c) Intelligent CAL

Fig 3 illustrates an Intelligent CAL system.

Here, the computer "contains" the Tutor together with an articulate Expert Model or Models (see Burton & Brown, 1979) and it is therefore the computer that constructs and tests models of the Learner, compares them with the Expert Models and teaching strategies in the Tutor, decides whether to intervene and if so, what to teach and how to teach it. Note though, that the Subject Matter Expert is still human in the system - as he must be. We must use him to represent the task domain from his point of view or perspective, but as we said earlier, the results of this process are not equivalent to the Subject Matter Expert. When we have completed this process we have represented the task domain and we have also represented a particular perspective on the task domain. We cannot seek to program the Subject Matter Expert, then, but rather to evolve a programmable representation of him (and the Learner and the teaching domain) upon which a machine based Tutor may operate effectively.

4. CONCLUSION

The purpose of this paper has been to present a rather simplified account of APU's philosophy and approach to research in the area of computer assisted learning. Three computer roles in learning systems were identified and described from a cybernetic point of view: these were Pre-CAL, CAL and Intelligent CAL. The Fighter Control Skills Trainer under development at APU is an example of a CAL system which will have the potential to improve considerably existing computer based

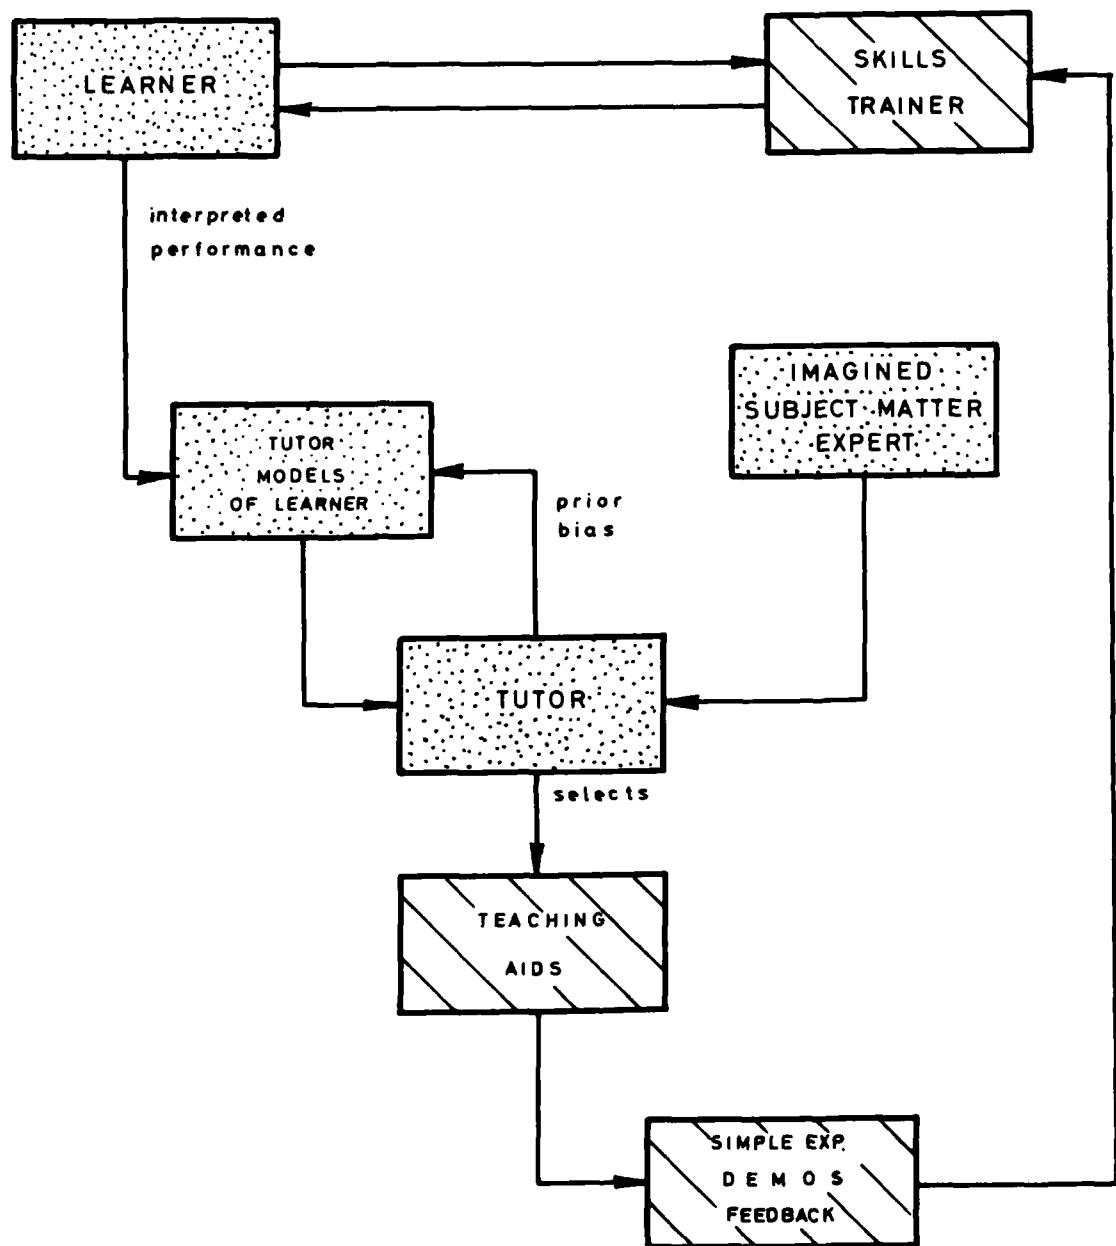
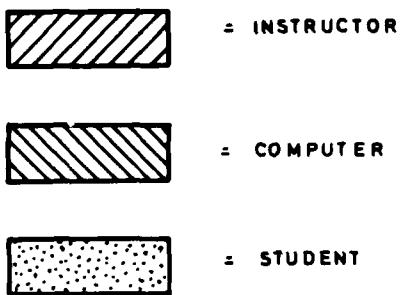


FIG. 2b CAL : STUDENT ONLY

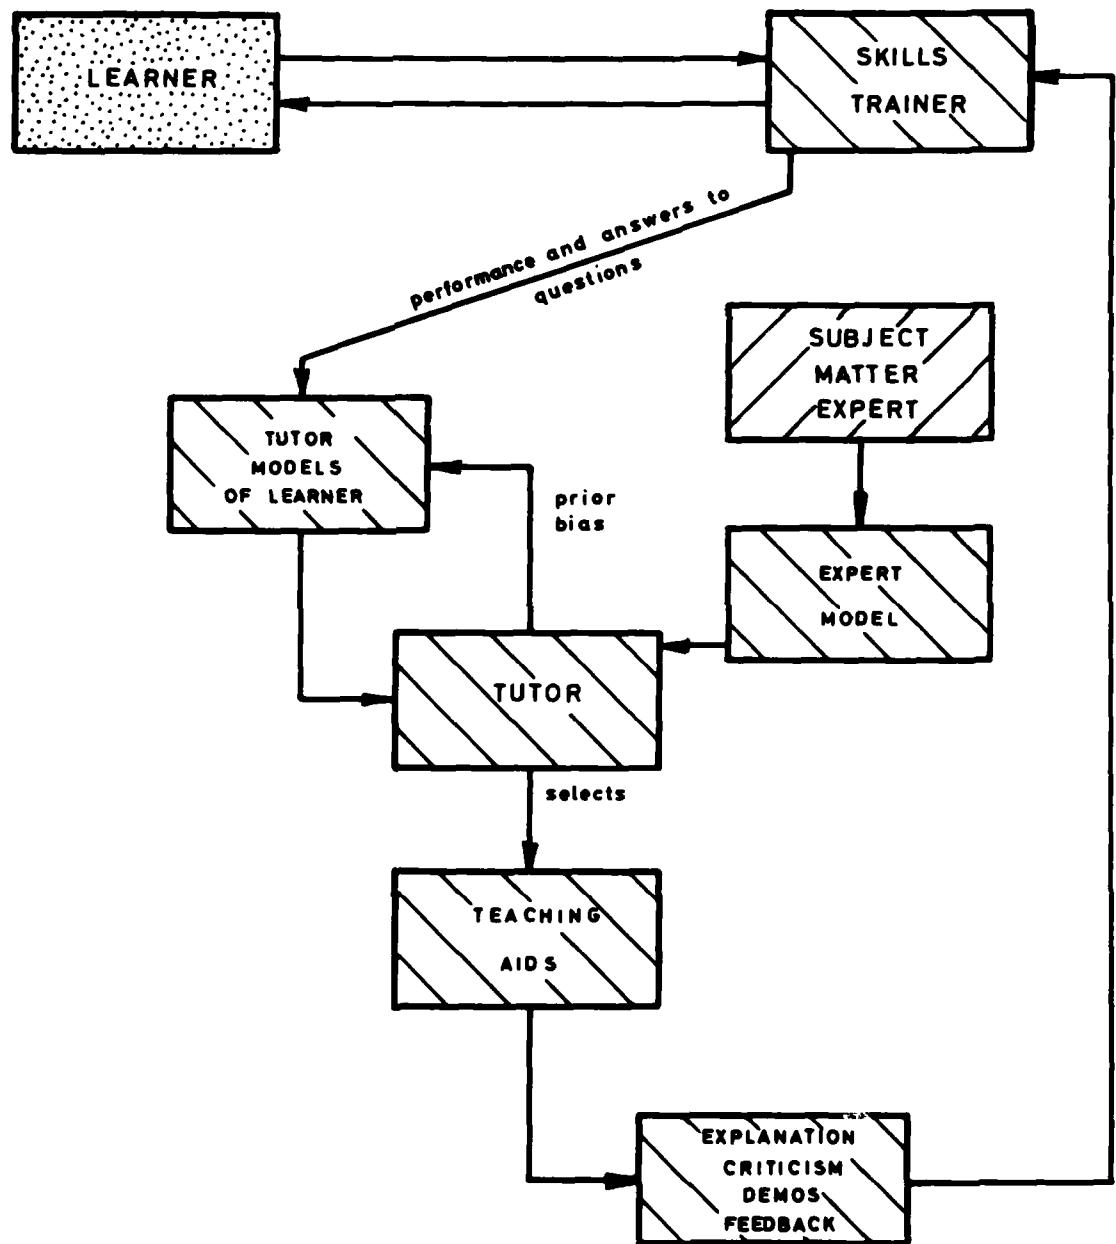
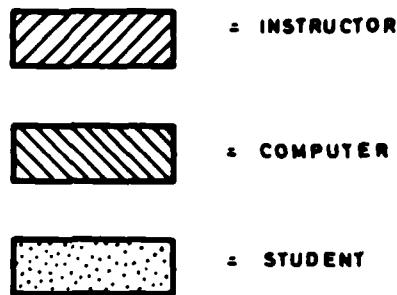


FIG. 3 INTELLIGENT CAL

training systems. Progress towards the even more powerful, intelligent systems is under way, and in particular, emphasis is being placed on the investigation of techniques for the representation of knowledge and user modelling.

REFERENCES

AUGSTEIN, S., & THOMAS L. Learning to learn. Brunel University. (1976)

BAINBRIDGE L., Verbal reports as evidence of the process operator's knowledge. Int. J. Man Machine Studies, 11, 411-426. (1979)

BURTON, R. R., & BROWN J. S. An investigation of computer coaching for informal learning activities. Int. J. Man Machine Studies 11, 5-24. (1979).

COURT, M. & BROOKING, B. Aircraft intercept control simulation studies: status report. AMTE (UNPUBLISHED) (1978)

FOLLEY J. D. Jr Development of an improved method of task analysis and beginnings of a theory of training. Port Washington, New York. USNTDC. NAVTRADEVCE 12 18-1. (1967a)

GAGNE, R. M. The conditions of learning. 2nd Ed Holt, Rinehart & Winston, London. (1960)

GOLDSTEIN, I. P. The genetic graph: A representation for the evolution of procedural knowledge. Int J. Man Machine Studies 11, 51-88. (1979)

GREGORY, R. Personalised task representation. Tech Memo AMTE(E) TM 79103. AMTE. (1979a) UNCLASSIFIED, UNLIMITED

GREGORY, R. Aircraft intercept control simulation studies: second status report. Tech Memo AMTE(E) TM 79105. AMTE. (1979b) UNCLASSIFIED

GREGORY, R. (in press) The air intercept control skills trainer. UNCLASSIFIED

MILLER, R. B. Task description and analysis. In R. Gagne (Ed) Psychological principles in system development, New York, Holt, Rinehart & Winston. (1962)

PASK, G. A Protolanguage (Lp: The THOUGHTSTICKER language). Systems Research Ltd., Woodville House, Richmond, Surrey. (1979)

PASK, G. (in press). Paper to be submitted to Int. J. Man machine relations.

SHEPPARD, C. (in press) Supportive operating systems, AMTE. UNPUBLISHED

SHEPPARD, C., GREGORY, R., & ROWLEY M. B. (in press): MISOS: The design of a mixed initiative supportive operating system. AMTE. UNCLASSIFIED

Reports quoted are not necessarily available to members of the public or to commercial organisations.

PAPER II: THE REPRESENTATION OF KNOWLEDGE: GOLDSTEIN'S GENETIC GRAPH

INTRODUCTION

The Genetic Graph (Goldstein, 1979) is a computer-based data structure that holds a knowledge representation. It exists as part of an intelligent Computer Aided Learning (CAL) system designed to teach the computer game of WUMPUS (see Appendix for a description of the game). The evolution of the Genetic Graph by Goldstein is important in at least two main ways. First, it is at the centre of the third version of a computer based WUMPUS ADVISOR (WUSOR III) and is one of the few existing systematic programmes to marry the technology of Artificial Intelligence (AI) with CAL.

Second, the Genetic Graph is a major attempt to change the emphasis of previous thinking with respect to intelligent CAL including WUSORS I and II: from how to represent the knowledge of the Subject Matter Expert to how to represent the evolution of expertise from the perspective of the Learner.

2. EXPERT BASED vs LEARNER BASED KNOWLEDGE REPRESENTATION AND MODELLING

Focus on the representation of Expert knowledge, rather than on the processes by which an individual acquires that knowledge had led to quite inadequate epistemological and pedagogical positions. These two positions are, in fact, two sides of the same coin. Epistemology is concerned with our beliefs about how Learners learn, while pedagogy is concerned with our beliefs about how Tutors should teach. Whatever epistemological assumptions one holds will lead directly - and may even dictate - a particular pedagogy.

To illustrate this point, let us look at what has happened with the Expert based approach to representation and modelling.

As Goldstein says, a fundamental assumption is that expertise consists of a set of facts or rules. The Learner's knowledge is modelled as a subset of this knowledge, and the Learner's learning as the process of expanding this subset until the subset is equivalent to the full Expert representation. This form of Learner modelling is referred to (Goldstein, op cit, Burton & Brown, 1979) as Overlay Modelling: the student's knowledge is overlaid on the Expert's to discover the difference. The Tutoring in such a system as this, "consists of encouraging the growth of this subset, generally by intervening in situations where a missing fact or rule is the critical ingredient needed to reach the correct answer". (Goldstein op cit).

Now, since the direction and activities of the Tutor are almost completely dependent on the epistemological assumptions inherent in the coaching system, it is to these assumptions that Goldstein has addressed himself; and the attempt to incorporate more elaborate and more powerful assumptions has led to the development of the Genetic Graph.

3. THE GENETIC GRAPH

At this point we may ask what these more elaborate and more powerful epistemological assumptions are: how have they been incorporated into the coaching system and how does this improve matters?

We should recall that the assumptions we are referring to are concerned with the processes by which the Learner acquires knowledge, and a key objective is stated early in Goldstein's (1979) paper: Goldstein wants to represent, via the Genetic Graph, the fashion in which new knowledge evolves from old by such processes as Analogy, Generalisation, Debugging and Refinement. The way he achieves this in the Genetic Graph is by representing Expert knowledge in terms of these processes. Goldstein's epistemological assumptions lead him, then, to incorporate these processes in his system of knowledge representation.

Let us consider some of these processes in more detail. Goldstein represents knowledge in the Genetic Graph at nodes. The nodes are connected by what he calls genetic links and these links reflect the epistemological processes that Goldstein has considered important to represent.

In Fig 1 are four rules (alias nodes, alias skills) linked by the three processes of Analogy, Generalisation and its inverse, Specialisation. There are three different specialisations of the same generalisation and these are analogous. The formal rule is: R' is analogous to R if there exists a mapping from the constants of R' to the constants of R . But there would seem to be a problem here: we will return to this in section 4, but it may be characterised as a question: Why is $R2.2$ not also analogous to its three specialisations?

The four nodes constitute a dense cluster providing the Tutor with multiple methods of explanation, ie one per link. Goldstein also argues that since these nodes are richly connected, then the Coaching System - and in particular, the Psychologist (see later) - can expect the Learner to acquire them more quickly and with less difficulty than nodes which are more isolated. These points anticipate a later part of the discussion but are mentioned here since they are, in fact, examples of pedagogical implications of the epistemological assumptions of Generalisation, Specialisation and Analogy.

In Fig 2 is added the process of Refinement and its inverse, Simplification. The formal definition is: R' is a refinement of R if R' manipulates a subset of the data manipulated by R on the basis of some specialised properties. But there is a problem here too which may also be characterised as a question: What is the difference, on this definition, between a refinement and a specialisation? This problem will be taken up in section 4.

The process of refinement is designed to represent the evolution of a rule to take account of a finer set of distinctions. There are five major phases of refinement in the WUMPUS syllabus and they correspond to five different Experts for the WUMPUS game. The process of refinement is used by Goldstein to get around a major criticism of Expert based systems. It is beyond reason, he argues, to assume that the

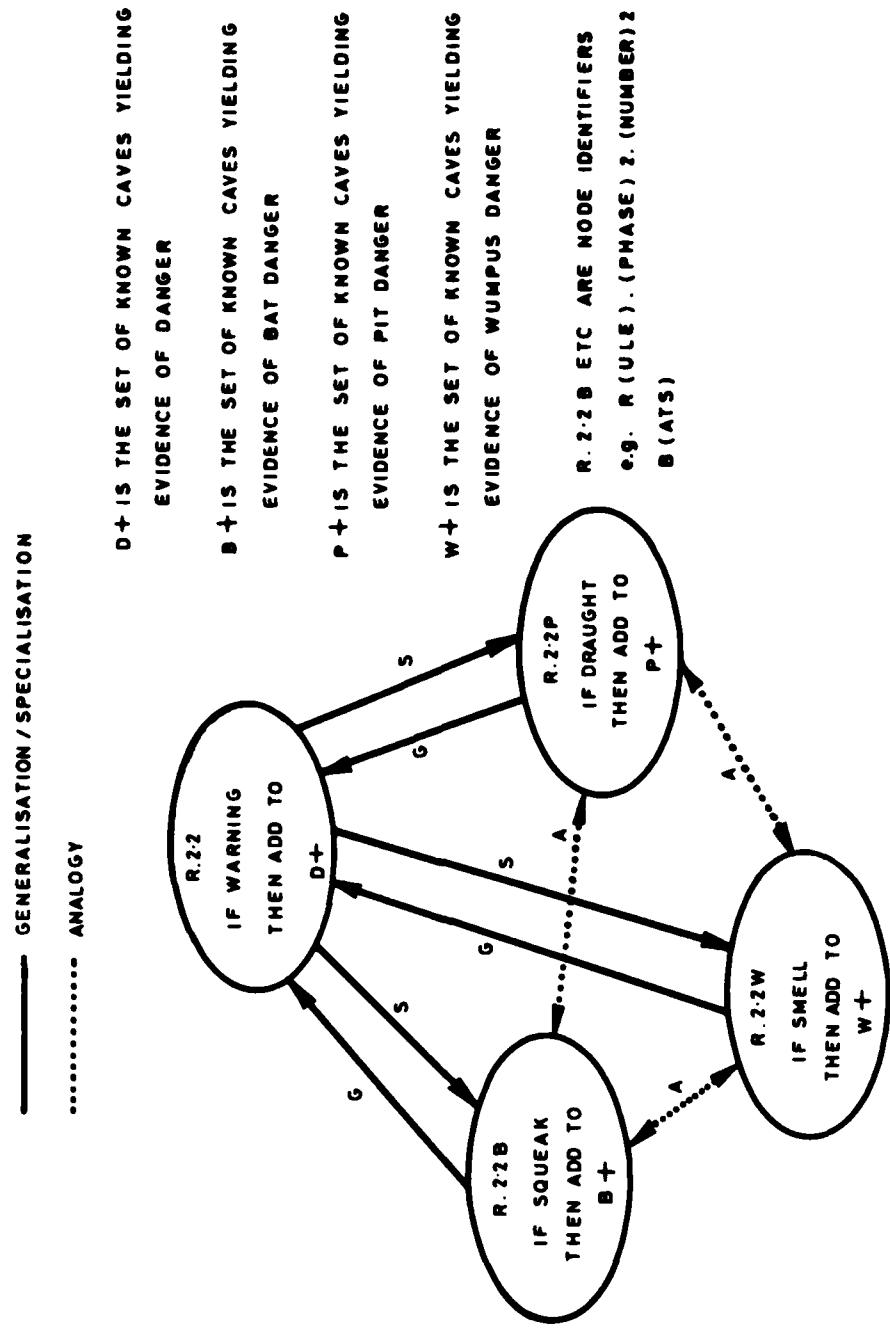


FIG. 1 SIMPLIFIED SUBSET OF THE GENETIC GRAPH FOR THE WUMPUS SYLLABUS

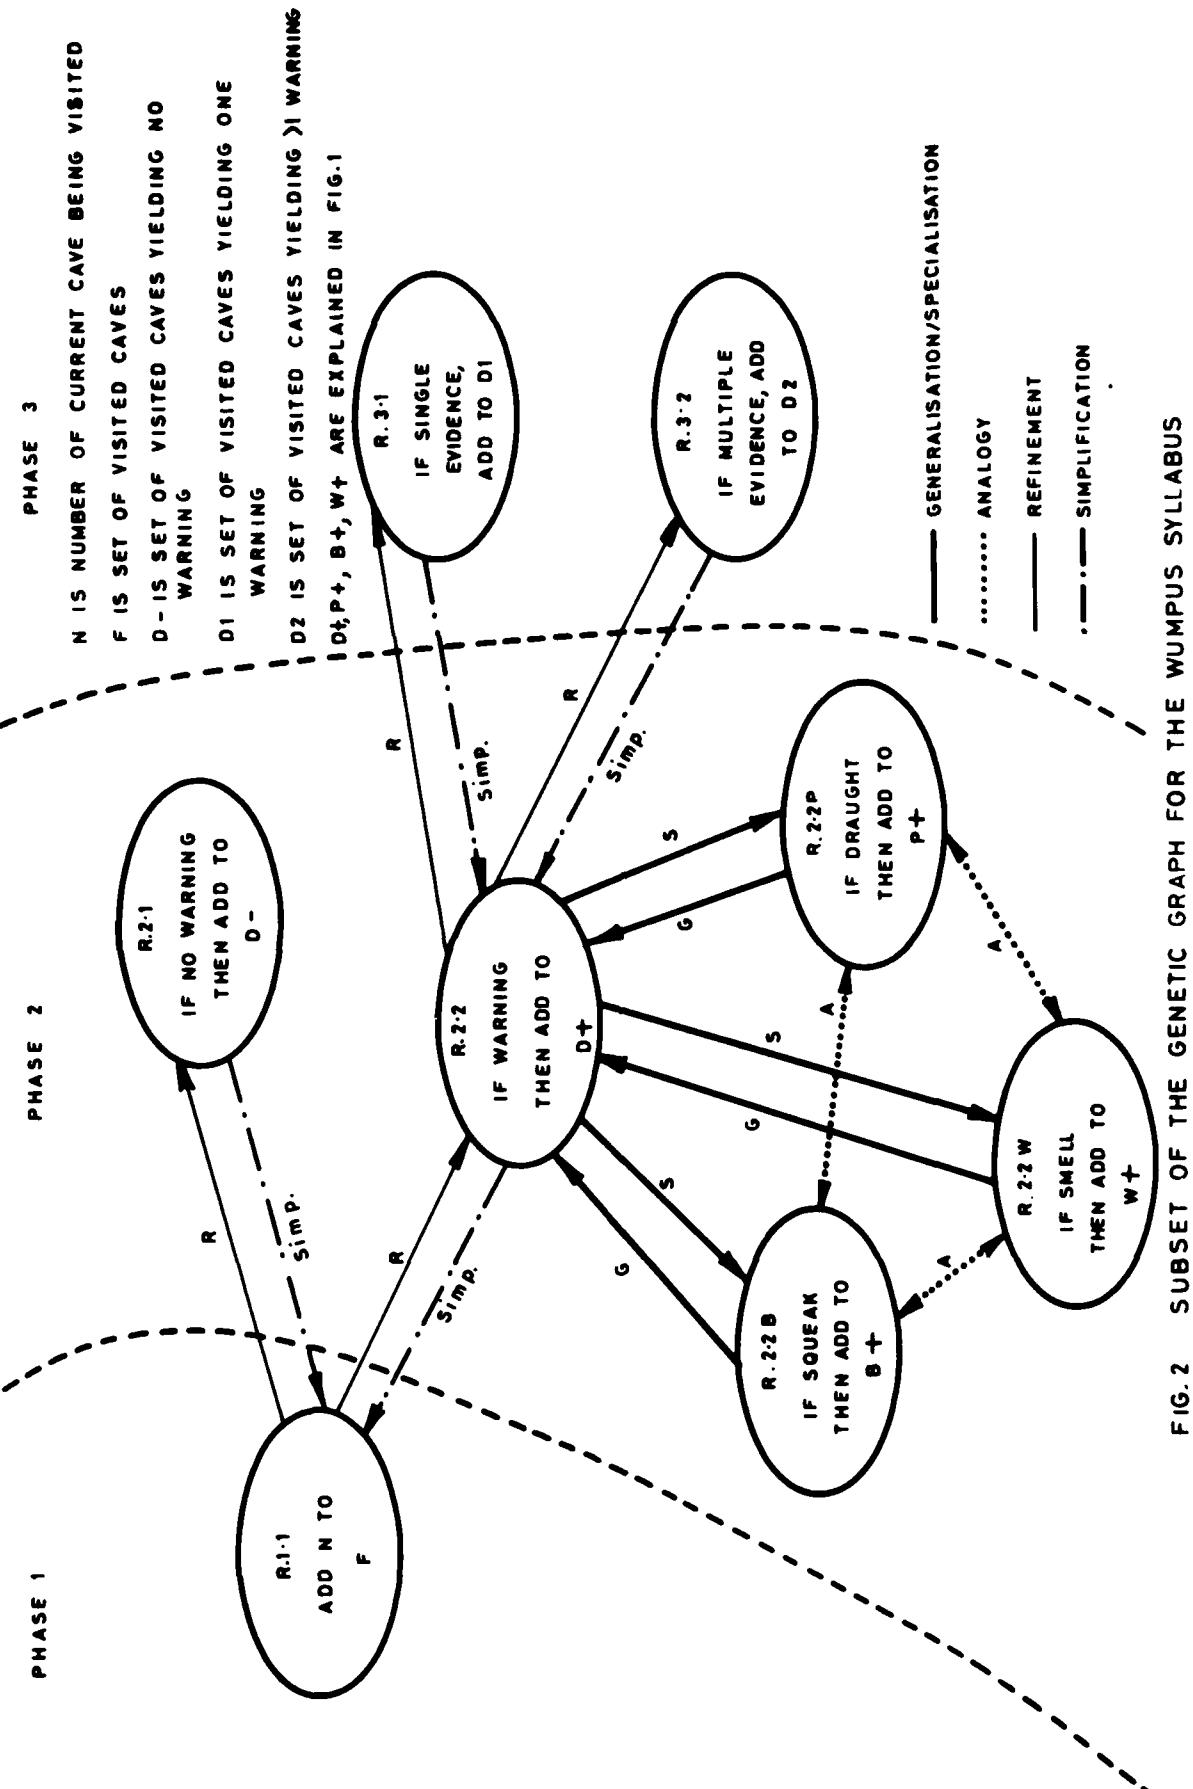


FIG. 2 SUBSET OF THE GENETIC GRAPH FOR THE WUMPUS SYLLABUS

student learns by elaborating a subset of the Expert's knowledge and further that this partial knowledge corresponds to some subset of the Expert's knowledge. Goldstein's answer, then, and indeed his epistemology is that the student learns by making ever finer distinctions in the rules he currently knows. Further, a student who has mastered Phase I cannot be said to possess a subset of Phase V, but Phase V itself in a grossly simplified form.

The Genetic Graph is then, a data structure for representing knowledge that incorporates Goldstein's assumption that certain processes underlie the way that Learners acquire knowledge.

The power of the Genetic Graph appears to be considerable. Once articulated (which poses another problem to which we will return), the data structure can be used to model a Learner in three different ways: the Learner's knowledge is modelled in terms of the nodes of the graph, his learning style in terms of the links, and his progress in terms of the paths in the graph.

Fig 3 shows the central role and Genetic Graph in Goldstein's WUSOR III system - the WUMPUS Coach.

With reference to Fig 3 we may consider the relationships between the Genetic Graph and the other components of WUSOR III.

Student Model

The Student Model is also a data structure holding what amounts to a personalised Genetic Graph. Based as it is on the Genetic Graph itself, the Student Model is able to provide the Tutor with three sorts of information:

- (a) What particular nodes (knowledge) the student has already acquired (K Model).
- (b) What particular sorts of links the student has used to arrive at his current knowledge frontier (L Model).
- (c) By what particular route the student has arrived at his current frontier. This information is important with respect to the learning complexity metric considered later.

The Student Model is an overlay not on the final skills of the Expert, then, but on the Genetic Graph which represents the evolution of those skills.

The Tutor

The Genetic Graph also supplies data to the Tutor, for the purposes of explanation. When the Tutor decides to break into the game with an explanation to the student, it must first decide what kind of explanation it is going to give. One of the ways it does this is by looking at the Student's L Model. If it sees, for example, that the Student's knowledge is connected with a preponderance of Analogy links, then it will be biased towards framing its explanation in terms of an analogy. Having

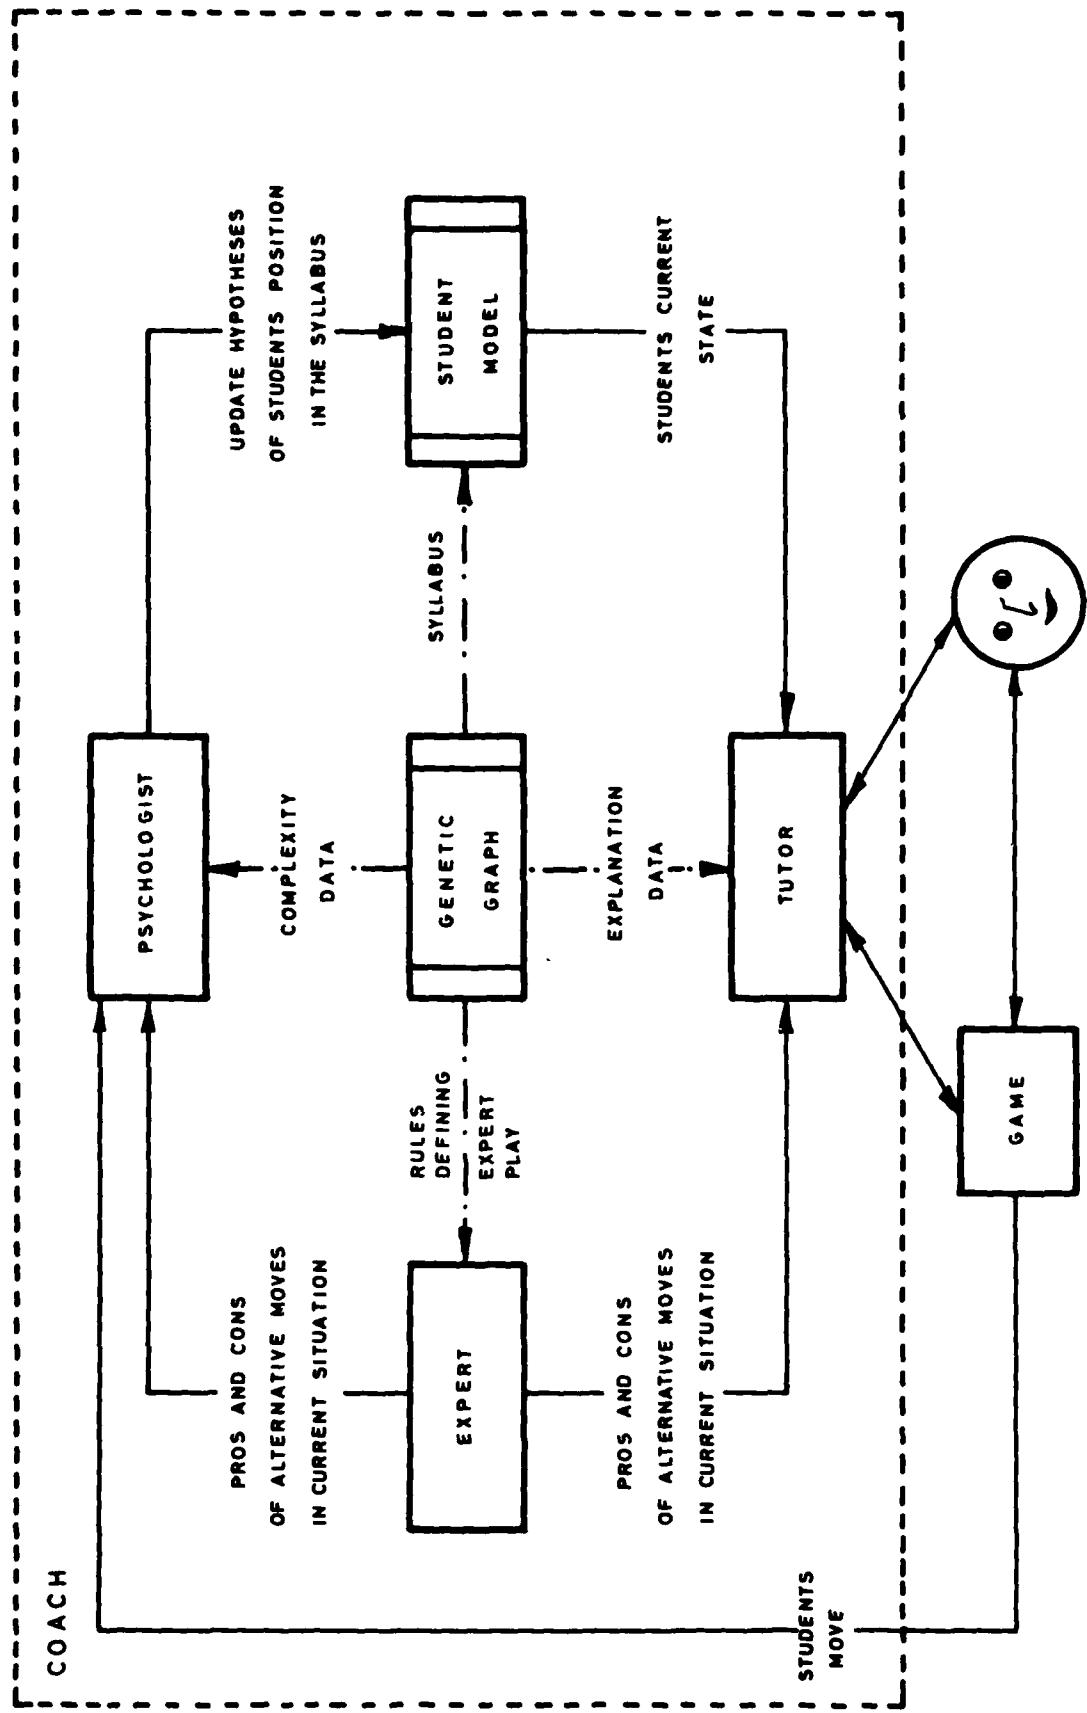


FIG. 3 GOLDSTEIN'S WUMPUS COACH: WUSOR III

made that decision, it then goes directly to the Genetic Graph for the appropriate detail, ie to find an appropriate analogy. As we have said before, the Genetic Graph can provide the Tutor with one type of rule explanation for every link that connects the to-be-learnt rule with a rule that the student has already acquired.

The Expert

In order to understand the importance of the Expert/Genetic Graph connection, it is necessary to consider the role of the Expert in the previous WUSOR II system (Carr & Goldstein 1977). In WUSOR II, the Expert performs three functions at each turn.

- (a) It generates all possible moves
- (b) It ranks them all
- (c) It identifies for each possible move including that selected by the student, what rules or skills are involved.

The assumption is that the player has learned all those rules involved in choosing his current move and rejecting its inferiors, and has yet to learn the rules needed to recognise superior moves. The Expert analysis in WUSOR II is used by the Psychologist to alter the Appropriate and Used columns for each rule. Now the point about the WUSOR II Expert is that if a rule is marked Appropriate, but not Used, the Psychologist is forced to the rather naive conclusion that the student does not possess that rule. With the reorganisation of WUSOR around the Genetic Graph (WUSOR III) however, the Psychologist receives not the Expert analysis, but five Expert analyses, corresponding to the five phases of the representation. The Psychologist's overall belief that the student possesses a given rule is a summation over the hypotheses of all five Expert players. The added constraint is that because the Genetic Graph incorporates an epistemology, the hypotheses generated by advanced "players" further and further away from the student's current frontier are assigned less and less weight.

What all this adds up to is that the belief metric of the Psychologist is greatly elaborated by the inclusion of the Genetic Graph via the Expert, and the modeller is thus made more sensitive to the student.

The Psychologist

Finally, in Fig 3, the Genetic Graph supplies complexity data to the Psychologist. Complexity is associated with graph density which was mentioned earlier. The less links connecting a particular node to the rest of the graph, then, clearly, the more isolated it is. The more isolated a node is, then the more the Psychologist is led to believe that it is difficult to learn.

The setting up of this belief ensures that the Psychologist will require more evidence that a Student has learnt the rule than in the case where the rule is part of a dense cluster.

Now, the difference between the learning complexity metric utilised by the Psychologist in the WUSOR II and in the WUSOR III systems is that in

WUSOR II, the Psychologist is made conservative in its belief that the student's behaviour has exhibited a particular rule when that rule is far from the frontier of the student's current knowledge state; in WUSOR III, however, the Psychologist is made to be conservative in believing that the student has acquired a particular rule when that rule is weakly linked to the student's knowledge frontier.

Complexity information is categorised as structural evidence because it describes the structure or topology of the Genetic Graph. The evidence supplied to the Psychologist by the Expert is implicit evidence because it infers what rules are required by different moves.

4. PROBLEMS

In the foregoing text, several problems have been briefly characterised and in this section we will take these up together with other more general comment. The discussion is organised by reference to three main problem areas and these relate firstly to Goldstein's glossing of what must have been major difficulties in physically producing a Genetic Graph representation of WUMPUS; secondly, to the possibility that the Genetic Graph is an artefact of the particular kind of domain that WUMPUS is, and thirdly to the view that Goldstein has replaced "Expert based" modelling with "Goldstein's epistemology based" modelling.

(i) The difficulties of Genetic Graph representation

The main question here is to ask how Goldstein has actually managed to produce a Genetic Graph for the WUMPUS domain? The problems must have been considerable and there is little account in Goldstein of how they were overcome. For example, how exactly was it decided when two rules were analogous and when they were not? How was it decided that multiple and single evidence rules were refinements of warning evidence, while draught, squeak etc were specialisations? The problems that Goldstein must have encountered while deciding how rules/nodes/skills related to each other are very similar to the problems Atkin reports with respect to deciding what level a particular Event in Q Analysis belongs to (personal communication) and certainly that Gregory has had in deciding what grain of detail an action belongs to in Personalised Task Representation (see Gregory 1979a). (There must be similar but, as yet, personally unencountered problems with respect to Pask's Condense operation in Lp as well - see Lp paper, this volume).

(ii) The Genetic Graph as an artefact of WUMPUS

Whilst other systems of knowledge representation embody some node reorganisation operation (Pask's condensation/expansion, Atkin's N^+ , H^+ and Gregory's Rolling In/Out operation) Goldstein only requires, it seems, the process of refinement spreading evenly out from Phase 1 through Phase 5. The fact that Goldstein can adequately represent the evolution of the domain knowledge of WUMPUS by the process of Phase refinement is, it is suggested, a result of the particular sort of domain that WUMPUS is.

In the first place, WUMPUS is not dynamic. The nature of the

game does not change as a result of some player action: the player's adversary (the WUMPUS) is neither purposive nor intelligent, and obstacles (pits and bats) do not change location or reside under external (purposive) control.

To play the game, the student does not have to generate tactics or a strategy as he would do in eg WEST (Burton & Brown 1979). Instead, he becomes more successful in direct proportion to his ability to make finer distinctions in the feedback that has always been available to him, albeit undetected. When a student can play WUMPUS as a Phase 5 Expert, it means he has an operational grasp of all the rules and their refinements, and thus at each move, for him, there is always only one move he can make. This is because the information fed back to him from his last move (or rather, his expert interpretation of it) actually dictates the move he should next make. For the sub-Phase 5 expert, the game moves are tactics that are obscured by unrecognised (undistinguished) feedback. The more expert one gets at WUMPUS, the more one's options are reduced. This may be contrasted with WEST or Command Decision Making in dynamic task environments where, the more expert one gets, the greater the possibilities are, because there are more strategies that one can be aware of.

(iii) "Goldstein's epistemology based" modelling

In devising the Genetic Graph, Goldstein's main objective was to move on from the Expert-based approach where the Student Model was viewed as a subset of the Expert Model. Instead of emphasising a learner-based approach where the Student Model captures the particular perspective that the student is taking, however, Goldstein appears to have incorporated his own interpretation of the learning process "from the learner's point of view" into a conventional Expert Model. This incorporation has resulted in the creation of five phases of expertise with the Student Model being overlaid on (expressed in terms of) these phases. This is not meant as a criticism and, indeed, Goldstein is most careful to say that he has done specifically this.

Whilst we at APU feel considerable support for Goldstein's rationale and reasons for overcoming the Expert based approach to modelling - and indeed see potential in the use of the Genetic Graph for representing certain Naval Tasks - for our general purposes it is not enough. We are in the process of examining the feasibility of genuine Learner-based modelling systems which in the context of a Supportive Operating System (Sheppard in press and Sheppard, Gregory & Rowley, in press) may be analysed and compared with (overlaid on) many Expert Models for subsequent appropriate Tutorial action. The central difference between Goldstein's Genetic Graph and our own enterprise is this: rather than use an epistemological connectivity, where the node relations are a theory of learning, we are implementing representational systems whose connectivities are independent of epistemology. Indeed, we conceive of the Tutor as consisting of different tutorial strategies (see Paper I, this volume) which are triggered by the results of the Learner Model/Expert Models comparison. (This will involve a meta teaching strategy to manage the Tutor's repertoire of strategies, making it

clear that there is always some stage at which one's own epistemology is manifested.) The important point, however, is that the Learner Model in our projected system will be genuine to the extent that the Learner is represented in his own terms rather than in terms of an arbitrary epistemology of an Expert Model subset.

5. CONCLUSIONS

The purpose of this paper has been to explore Goldstein's Genetic Graph system for representing knowledge as an advance on Expert based modelling. That it is an advance is without question and the primary mechanism of the improvement derives from Goldstein's concern to represent not just Expert knowledge but also the evolution of Expert knowledge from the perspective of the Learner.

To achieve this, Goldstein has established an epistemology as the connective structure of the Graph. Thus the links between nodes themselves represent "learning processes" such as Generalisation, Analogy and Refinement. As a result, the Genetic Graph allows a much more powerful modelling facility (ie the Coaching system is more sensitised to development of the Learner's expertise) this being achieved (to state it simply) by the Graph's capacity to yield information about the Learner's knowledge state, his learning style and his route through the syllabus.

The Genetic Graph is not, however, a true learner based modelling facility such as we are pursuing at APU. (To be sure, it must be emphasised that Goldstein has not made such a claim.) Further, it could be that the Genetic Graph is not domain-type free either, with more complex kinds of connectivity being required for dynamic task environments in which the user's adversary is as intelligent and purposive as he.

Despite these comments, the Genetic Graph may be of considerable use for representing one or two Naval tasks on which APU staff are working and, in any event, the exploration of Goldstein's system has been of considerable use to us in helping as to formulate our own programme of research on intelligent Supportive Operating Systems.

REFERENCES

BURTON, R. R. & BROWN, J. S. An Investigation of Computer Coaching for Informal Learning Activities. Int J Man Machine Studies 11, 5-24. 1979.

CARR, B. & GOLDSTEIN, I. P. Overlays: A Theory of Modelling for Computer Aided Instruction. Mass. Ins. Tech AI Lab. AI Memo 406. AD-A052 210. 1977.

GOLDSTEIN, I. P. The Genetic Graph: A Representation for the Evolution of Procedural Knowledge. Int J Man Machine Studies 11, 51-88. (1979)

GREGORY, R. Personalised Task Representation. Tech Memo AMTE(E) TM 79103. AMTE UNCLASSIFIED, UNLIMITED.

SHEPPARD, C. (in press) Supportive Operating Systems. AMTE UNPUBLISHED

SHEPPARD, C., GREGORY, R. & ROWLEY, M. B. (in press) MISOS: The Design of a Mixed Initiative Supportive Operating System. AMTE UNCLASSIFIED.

Reports quoted are not necessarily available to members of the public or to commercial organisations.

APPENDIX

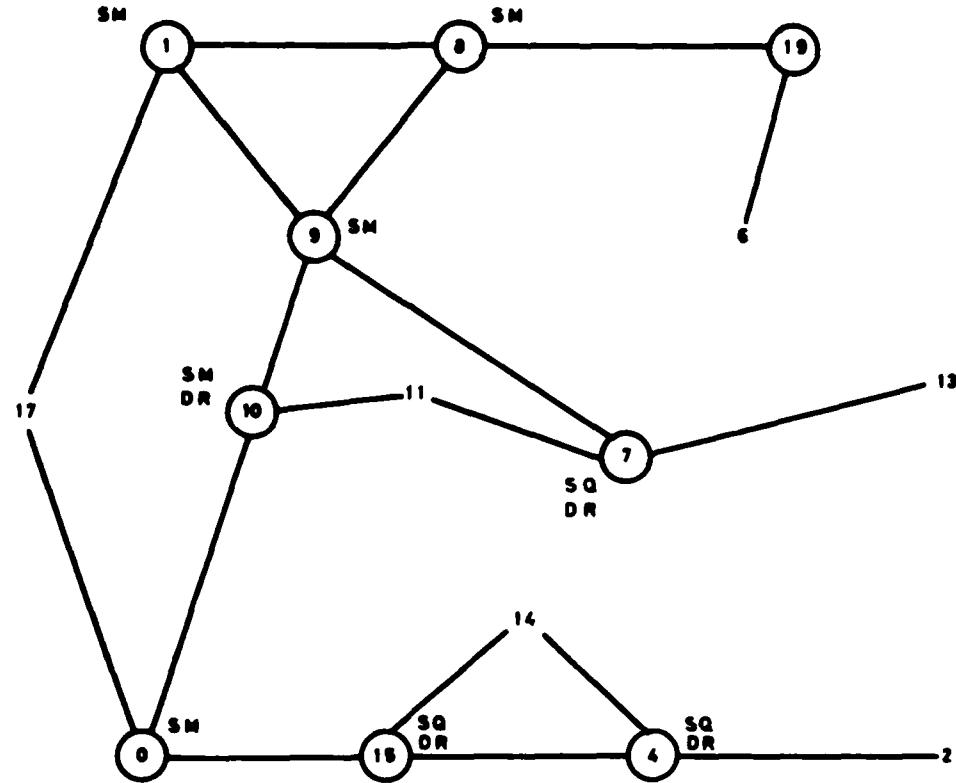
WUMPUS, an Intellectual Game*

"The WUMPUS game was invented by Gregory Yob (1975) and exercises basic knowledge of logic, probability, decision analysis and geometry. Players ranging from children to adults find it enjoyable. The game is a modern day version of Theseus and the Minotaur. The player is initially placed somewhere in a randomly connected warren of caves and told the neighbours of his current location. His goal is to locate the horrid WUMPUS and slay it with an arrow. Each move to a neighbouring cave yields information regarding that cave's neighbours. The difficulty in choosing a move arises from the existence of dangers in the warren: bats, pits and the WUMPUS itself. If the player moves into the WUMPUS' lair, he is eaten. If he walks into a pit, he falls to his death. Bats pick the player up and randomly drop him elsewhere in the warren.

But the player can minimise risk and locate the WUMPUS by making the proper logicistic and probabilistic inferences from warnings he is given. These warnings are provided whenever the player is in the vicinity of a danger. The WUMPUS can be smelled within one or two caves. The squeak of bats can be heard one cave away and the draught of a pit felt one cave away. The game is won by shooting an arrow into the WUMPUS' lair. If the player exhausts his set of five arrows without hitting the creature, the game is lost."

Fig A illustrates a typical intermediate state a player might reach.

* Extracted from Carr and Goldstein (1977)



**CIRCLED CAVES HAVE BEEN
VISITED BY PLAYER**

NUMBERS ARE CAVE IDENTIFIERS

SM - WUMPUS WARNING

SQ - BAT WARNING

DR - PIT WARNING

FIG. A AN INTERMEDIATE STATE IN A TYPICAL WUMPUS GAME

PAPER III: THE REPRESENTATION OF KNOWLEDGE: PASK'S L_p

INTRODUCTION

The purpose of this paper is to provide a generalised exposition of the language of knowledge representation, L_p, which has been at the centre of Gordon Pask's research for the past few years. It is not the intention to trace the evolution of L_p since this is done admirably elsewhere (Entwhistle, 1978, Pask, 1975, 1976, 1979, in press). Rather, there is a requirement to explain what L_p is and what it does, the latter by reference to a particular implementation of L_p called THOUGHTSTICKER.

L_p is a protolanguage; that is, a language which is fundamental but rudimentary. This may be put another way: " 'Language' users may employ 'language' to communicate with each other, but primarily, they modulate some refinement of L_p". (Pask, in press). L_p is used to model the conversational process between two entities; that is two entities who are sharing concepts or understanding each other. L_p is, crucially, active, a process (or collection of processes) that represents the concepts, memories, understandings, learning, agreements, information transfers, etc in intelligent systems (including human brains) as processes. This is quite distinct from the traditional view of concepts, memories etc as states, a distinction to which we shall return. The conversing entities need not correspond to two human brains: conversations may (do) occur in one brain, eg when a person considers the relationship between, say driving a car and riding a motorcycle. Equally, conversations may occur between systems that do not involve any biological material. Indeed L_p is designed to represent the process of understanding over an arbitrary collection of arbitrary types of processor. L_p is the language of Conversation Theory (Pask 1976). The observables of Conversation Theory are agreements, and the sharp valued observable events are agreements between participants over an understanding. Here, understanding has a technical meaning. To understand a concept, one must demonstrate the concept by executing some procedure that realises it. In the THOUGHTSTICKER implementation of L_p, knowledge is represented at uniquely named nodes designating particular concepts or topics which by their interconnectivity comprise a mesh. There are two main ways in which topics may be connected. First, every topic is uniquely identified by its derivation from other topics. Knowledge about each topic thus entails knowledge about particular, other topics, and hence the mesh is known as an entailment mesh. Second, topics and their derivations may be analogically connected which in the simplest case, constitutes an isomorphism between distinct universes of discourse.

The main body of this paper is organised around a systematic explanation of some of the main concepts and operations that constitute L_p. Particular areas of coverage include Derivation, Coherency and Organisational Closure; Rule of Genoa and Analogy; Pruning and Selective Pruning; Saturation; Condensation and Analogy; and Isomorphic Inference. At the same time, this 'coverage' should be regarded as a superficial treatment of the richness, subtlety and power of L_p. As

already indicated, the paper is intended as an introduction to Pask's and others' papers to which the interested reader is directed in the references for more comprehensive and detailed treatment in L_p .

2. L_p : OPERATIONS AND OPERATIONAL CONCEPTS

2.1 At the centre of L_p and Conversation Theory is the proposition that no piece of knowledge or concept is isolated in the brain of the knower. Related to this is the further proposition that knowing something is an act or process rather than a mental state. It is an unfortunate constraint of the present medium that whilst L_p must be regarded as a dynamic process which models other processes (concepts, memories etc) it may only be described here by means of static inscriptions and diagrams in two dimensions. All that can be done at this juncture is to recommend the reader to be constantly aware that the diagrams we will use are static representations of essentially dynamic processes.

Let us examine a now famous example (Pask 1979). Imagine that one person (Adam) wishes to converse with another person (Eve) about what his concept of circle is. For example, Fig 1 shows one way in which Adam may represent how he knows what a circle is in L_p .

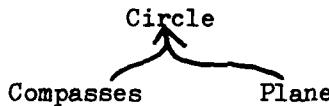


Fig 1 Adam's representation of a circle

What Adam means here is that he can regenerate his concept of circle by mentally combining his idea of a pair of compasses and a plane. But what L_p requires is that if this is to be upheld as a valid derivation of Circle, then the other two possible permutations must also make sense to Adam viz:

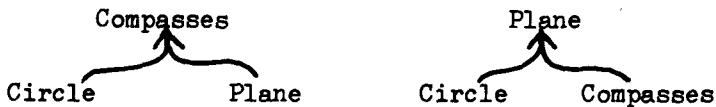


Fig 2 The representation is cyclic

In fact, Adam deems it possible, for him, that if he knew what Circle was and what Plane was, then he could get an idea of a device for inscribing a Circle on a Plane. A similar justification applies to the third case.

Returning to Fig 1 we may now conclude that Adam's concept of Circle is stable or cyclic or coherent (as are also the concepts of Compasses and Plane when the group or bundle is viewed from those perspectives, Fig 2). So far, in Fig 1, the concept Circle is said to have one kernel denoted by the derivational arrow head. If in the completed mesh (of the conversational domain, say

"Geometric Figures"), the situation remains as it is, then Circle is said to have a collective derivation. But in Adam's case, he has several other ways of thinking about (ie deriving) the concept Circle. In this case the concept is said to have a distributive derivation. Fig 3 shows two more which Adam instates as coherent bundles (ie cyclic).

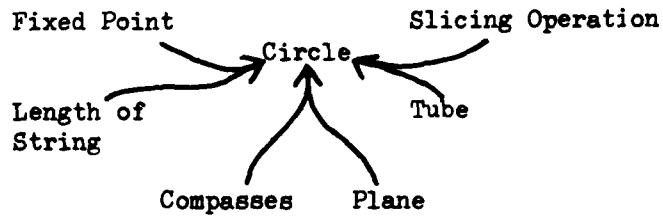


Fig 3 Three coherent bundles

Let us now consider Adam's task of helping Eve to understand what he means by Circle. Adam establishes that Eve understands (in its technical sense) what a Tube is, and further what it is to Slice something, Adam then asks Eve to put these together in her head and hence get an idea or image of what a Circle must be.

It will be noticed that the derivations of Circle are not simple breakdowns of the concept: the structure is not taxonomic. In fact the concept Circle is Adam's repertoire of the ways in which it can be understood. Further, for each of these ways, the coherency requirement ensures that all the permutations obtain at the same time. This is another way of saying that each topic (eg, concept) in the coherent bundle is recalled or derived from a combination of the other members. Concepts are remembered in terms of other concepts. The coherency requirement therefore embraces the idea of memories, concepts etc as processes. The memory of Circle is the process of deriving it from at least two other concepts. But there is a further point to be made about the coherency requirements: because Lp uses not a true/false logic but a logic of coherence, Adam is not being made to assert in Fig 3 that this is what a circle actually is; rather that this is how he consistently and reliably thinks about Circles. Fig 3 defines his concept of Circle rather than the concept of Circle in any 'absolute' terms. It is a basic tenet of Conversation Theory (which Lp has been designed to express) that a concept belongs to some organisation, eg Adam. A concept is regarded as a stable repertoire of ways of knowing about that concept. Another basic tenet is that a concept is also a procedure. Upon execution, a concept is manifest as a process. In this respect Adam's concepts are like skills. The execution of concepts (skills) may give rise to a description or to a behaviour, or both. In this way, Lp representations of knowledge in THOUGHTSTICKER are both descriptive and prescriptive. It is the novel treatment of concepts, memories etc as coherent processes that allows knowledge to be represented both descriptively and prescriptively simultaneously. This treatment lends considerable power to Lp.

Let us try to represent Eve's construction of a description of

a Circle (C) from Adam's invitation to Eve to consider Compasses (Co) and a Plane (P) which are concepts that Eve already understands.

Co_E, P_E

are descriptions of Compasses and Plane held by Eve.

(i) $\underline{\text{Ex}} \text{DB}(Co_E, P_E) \Rightarrow C_E$ The execution of a Description Building Operation enables E to combine Co with P and so produce a description or image of C.

Now concepts that are understood (in its technical sense) are both descriptions and procedures ie to understand a concept, it is not enough to claim to understand it, but, as we have said, one must demonstrate one's understanding by creating an example in some way. Another way of expressing this difference is in terms of theoretical knowledge vs operational knowledge. At the moment Eve has procedures and descriptions for Co and P, but only a single image of C. To complete her understanding of C, Eve must execute a Procedure Building Operation:

↓

(ii) $\underline{\text{Ex}} \text{PB}(Proc_E \text{Co}, Proc_E \text{P}, C_E) \Rightarrow Proc_E \text{C}$

C_E (arrowed) is the description of C by Eve produced by the DB operation (i). It is inserted in the PB statement to derive a procedure for producing an example of C. Having produced this procedure, $Proc_E \text{C}$, it now becomes part of Eve's concept of C. But we have not finished yet. Now that a demonstrable understanding of Circle exists in Eve (in terms of Compasses and Plane), a new understanding of Compasses can be brought about in terms of the other two. The same is true for Plane.

(iii) $\underline{\text{Ex}} \text{DB}(C_E, P_E) \Rightarrow Co_E$ This is a new image of Compasses because C_E did not exist before.

and (iv) $\underline{\text{Ex}} \text{PB}(Proc_E \text{C}, Proc_E \text{P}, Co_E) \Rightarrow Proc_E \text{Co}$ This is a new procedure for creating (in some way) Compasses which is added to E's concept repertoire of Compasses.

similarly

(v) $\underline{\text{Ex}} \text{DB}(C_E, Co_E) \Rightarrow P_E$ This is a new image of Plane

(vi) $\underline{\text{Ex}} \text{PB}(Proc_E \text{C}, Proc_E \text{Co}, P_E) \Rightarrow Proc_E \text{P}$ This is a new procedure for creating a Plane which is added to E's concept repertoire of Plane.

It is in this way that knowledge is coherently represented in L_p both prescriptively and descriptively. We do not want to know

just how knowledge relates to the rest of the world, but what those relations allow us to do. The coherency requirements both of Fig 1 etc and of the DB and PB operations mean that the process of memories or knowledge is stabilised in Eve's brain. This stability or coherence (eg of Circle, Compasses, Plane) is also known as organisational closure. If a system (eg part of a mesh, or mesh or Adam), is organisationally closed then it acquires an integrity or an autonomy. But the same system is not informationally closed. The concepts in the organisationally closed structure in Fig 3 may be accessed by any other concepts introduced into the mesh and this access may be either derivational or analogical in character. Whichever it is, it does not affect that system's autonomy. Circle in Fig 3 is part of three organisationally closed bundles, and informationally open to all three. Likewise, Adam, though organisationally closed, is informationally open: Eve can ask him things and Adam can answer. Specifically, Eve has a conversation with Adam, the observables being the agreements (including agreements to disagree) that emerge between them. In fact, any agreement that comes about between Adam and Eve is represented in L_p as an analogy, whereby for the purposes of their conversation Adam and Eve become analogically related. The distinction between them must be maintained (they are two different organisationally closed systems) but if they are to share each others concepts, then those concepts must be coherent with each other.

It is not enough that L_p should use a logic of coherence. If two otherwise independent, organisationally closed (coherent) systems (eg parts of a mesh, different meshes, Adam and Eve, computing media) are to converse (share concepts), then the language that represents this conversation must use a logic of coherence and distinction. The mechanism in L_p that detects the need for a distinction is the Rule of Genoa, and the relation that links two distinct entities in order that they may be informationally open is Analogy. These two mechanisms are dealt with next.

2.2 The Rule of Genoa and Analogy

We may recall from the Introduction that the topics in an entailment mesh are required to be uniquely identified by their derivations from other topics. If this uniqueness is compromised, ambiguity is introduced into the mesh which must be resolved. Consider Fig 4.

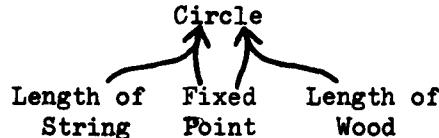


Fig 4 The contravention of Genoa

Here, Adam is trying to add derivations of his concept of Circle to make the connectivity of his mesh richer. He has been careful to use different topic names and is satisfied that he has

created two coherent bundles, each of three concepts. On the face of it, the structure looks perfectly reasonable. But let us work through the implications. To do this we will use a slightly different notation: $C(S, FP)$ simply means that Circle is derivable from String and Fixed Point. Thinking about this bundle first, we obtain:

- (a) $C(S, FP)$
- (b) $S(C, FP)$
- (c) $FP(C, S)$

For the second bundle, we obtain:

- (d) $C(FP, W)$
- (e) $FP(C, W)$
- (f) $W(C, FP)$

If we now examine (b) and (f), we find that the same derivation (C, FP) gives rise to two different concepts, S and W. This contravenes the Rule of Genoa (pointed out to Pask by Vittorio Midoro of Genoa) that differently named concepts may not have the same derivation. It should be obvious why this is important. The contravention of Genoa is tantamount to asserting that two differently named concepts are identical. There are a number of ways that Adam may resolve the problem. Adam reflects that although String and Wood are different concepts, they are both useful here, in conjunction with Fixed Point because they share the same important characteristic, ie they both have the capacity for tautness or rigidity.

In this way, the two concepts are analogous. They have differences (eg in origin, in texture etc), but are similar for the purpose of generating the concept of Circle. Adam's reflection focuses for him how he is utilising his concepts of String and Wood, but he realises that to assert an analogy relation between the two, although valid, does not solve Genoa. He would still be left with the original problem of the same derivational set (Circle and Fixed Point) contributing to more than one node.

Adam considers doing away with String and Wood altogether and replacing them with a single node which he labels Rigid Length of Unspecified Material. In doing this, Adam submits himself to a Condense operation which automatically moves him out of this level of representation into a mesh comprising higher order concepts (see section 2.5). The Condensing of analogies is a vital operation in Lp, for it reveals precisely what one's higher order concepts are, and we will return to it later (section 2.5). Since Adam's reflections have allowed him to appreciate the analogous relation between the ways he is using his concepts of String and Wood, he asserts the relation in Fig 5. To be clear this analogy is not part of the resolution of Genoa, but rather a move that Adam generates because of his insight.

The way Adam does resolve Genoa, thus remaining at his current level of representation, is expressed by the analogous relation

Adam has agreed to between Circle' and Circle" in Fig 5.

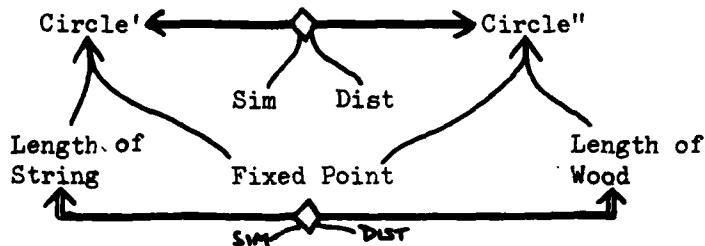


Fig 5 The resolution of Genoa

Recalling that concept derivations are not only descriptions of the concept, but prescriptions for realising an example of the concept, Adam realises that he cannot produce the same Circle with $S + FP$ and with $FP + W$. He therefore sees that he must distinguish between Circle' and Circle" though he sees them both as having similarities (they are both Circles). Adam resolves Genoa by asserting an analogy relation between Circle' and its derivation, and Circle" and its derivation, characterised by similarities (Sim) between these two views of Circle and distinctions (Dist). He need not state what these similarities and distinctions are if he does not want to, although he must if he submits to a Condense operation (section 2.5). When two concepts must be made distinct in this way, the analogy relation is used to bridge the gap. When a mesh is viewed as a surface on which concepts and their relations are inscribed, an analogy may be seen as a repair to a torn topology; Pask refers to this view of an analogy as an 'essential singularity'. But the mesh may also be seen as a process; in this view of things, an analogy represents a process bifurcation. The process of knowing about the concept of circle splits into distinct parts of universes of discourse. It may be helpful to think of process bifurcation as a "cell division of ideas" * - a kind of asexual reproduction. This biological metaphor may be extended further, for the condensing of an analogy to a higher order mesh may be likened to sexual reproduction in which two entities liaise and give rise to a single individual belonging to the next generation. This second kind of reproduction is taken up in section 2.5, when we come to consider the analogy as representing the process of agreement and hence the process of information transfer between intelligent systems.

2.3 Pruning, selective Pruning and Superimposition

Let us consider Fig 3 once more, but this time add the cyclicity relations required by coherency, Fig 6.

* I owe this metaphor to Paul Pangaro

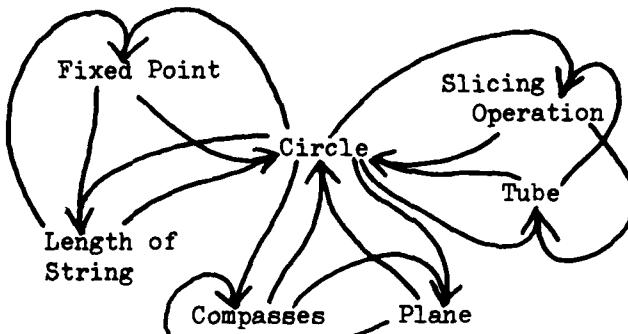


Fig 6 A simple mesh

To Prune the mesh we designate a particular node as the head node and unfold the mesh underneath it. In this way we examine the mesh from the particular perspective afforded by our selected head node. Straightaway we may see that Fig 3 is not a mesh, but a Pruning of the mesh in Fig 6. Let us Prune Fig 6 under Tube. Fig 7 gives the result.

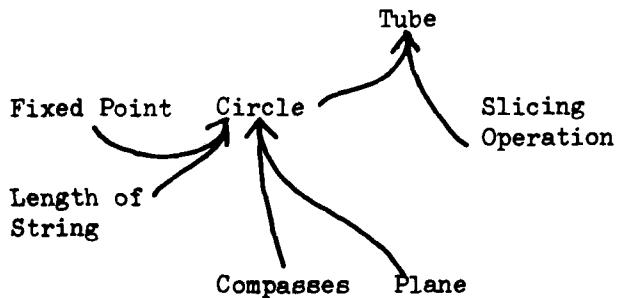


Fig 7 A pruning of the mesh in Fig 6

In its true state with all cyclicity added, an entailment mesh is a mass of nodes and their interconnections, which exists and resonates but has no purpose. To allow us to use a mesh for a particular purpose, we can unfold it with the Pruning operation. If, from Fig 6, Adam wishes to converse with Eve, not about Circles but Tubes, Adam selects Tube as the head node or conversational goal and executes Prune. Adam wants to explain to Eve his concept of Tube and has just one way of doing so (in this mesh). He thus asks Eve to demonstrate her understanding of Circle and Slicing Operation so that he can then ask Eve to combine the two to get an image of Tube. Eve knows what a Slicing Operation is, but fails to satisfy Adam that she knows what a Circle is.

Since Adam only has one way of communicating with Eve about Tube and this necessarily involves knowing about Circles, Adam realises he must first get Eve to acquire a concept of Circle. Adam has two ways of doing this, and chooses the Fixed Point/String derivation since Eve demonstrates that she has acquired both of these concepts. And so the conversation goes on.

This is a rather superficial treatment but illustrates the points that Pruning under a specified goal node reveals not only

a perspective on the mesh, but a plan of action. More accurately, Pruning a mesh reveals all the possible plans of action for a particular goal. Selective Pruning (Selprune) reveals just one. There are two possible Selective Prunings under head node Tube for Adam. They are given in Fig 8a and 8b.

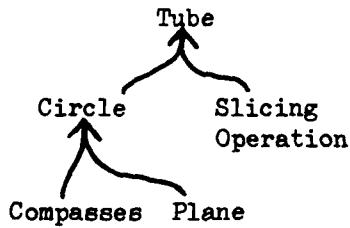


Fig 8a Selprune 1

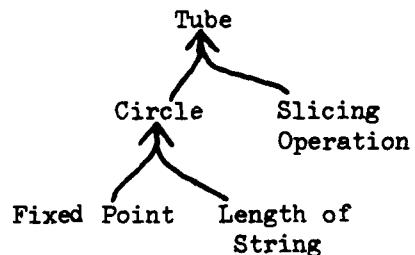


Fig 8b Selprune 2

It should be clear that a Pruning of a mesh will include every node in the mesh, but unfolded from a specific perspective. Superimposition is an L_p operation that reconstructs the mesh by combining or superimposing all possible mesh Prunings. Similarly, superimposition of all Selective Prunings under a particular node reconstructs the Pruning under that node.

2.4 Saturation

It is a feature of intelligent entities that they attempt to make connections before they are formally stated. This is what thinking is. The more interconnected a mesh (or brain) is, the more stable it is and the more efficient it is. As Frank Lloyd Wright put it "An expert is a man who has stopped thinking - he knows". In L_p , Saturation is an autonomous process that attempts to add as many kernels as possible for each topic until and before bifurcation occurs. (See Clark, 1979).

Let us again consider Fig 6 (the mesh, not the Pruning). The way that Saturation would work on this would be to suggest to Adam that, for example, Circle could be derived from Plane and Tube, or Sl . Operation. It is up to Adam to consider and respond to these suggestions, agreeing with, or rejecting them. When, for any mesh, no more combinations can be agreed, then the mesh is said to be Saturated. It is interesting to note that fully Saturated organisationally closed systems occur at the boundary of bifurcation. One more demarcation from within the system would contravene Genoa and thus, demolish it, calling for analogical repair.

There is a distinction to be made between Saturated and fully Saturated systems. We can see this by referring to the correspondence between L_p 's Saturation process and entities known as Steiner Rings, Fig 9.

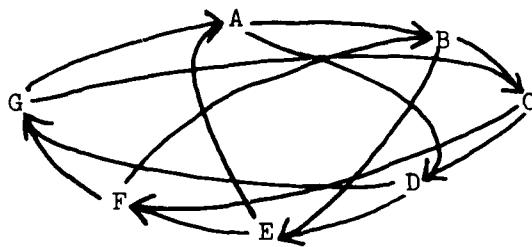


Fig 9 A Steiner Ring

Assuming that we are concerned with bundles of three nodes, the first number of nodes that can exhibit full Saturation is, of course, the minimal case, 3 (every node is derivable from the other two). Steiner showed that the next number of nodes exhibiting full Saturation (alias complete stability) was 7, (drawn in Fig 9). The next number after that is 9, then 13, 15, 19, 21, 25 and so on. This means that for the mesh in Fig 6 (with 7 nodes) it is theoretically possible to reach full Saturation. (This theoretical proposition is modified, of course, by Adam's willingness to agree that all combinations make sense to him). Meshes whose total number of nodes does not correspond with one of Steiner's numbers can never be, even theoretically, fully Saturated. But there is a further point here: adding kernels to a mesh in different orders results in different Saturation possibilities and this is partly because the kernels may consist of varying numbers of nodes. This has implications for training systems where students who acquire knowledge in a way which coincides with optimal Saturation processes will construct a more stable knowledge structure than those whose knowledge acquisition is at variance with those processes. At any rate, the former student-type should learn faster. This has relevance to courses in learning to learn.

2.5 Condense & Analogy

Let us backtrack a little and consider once again Adam's resolution of the Genoa problem in Fig 5. The figure is redrawn below for easy reference.

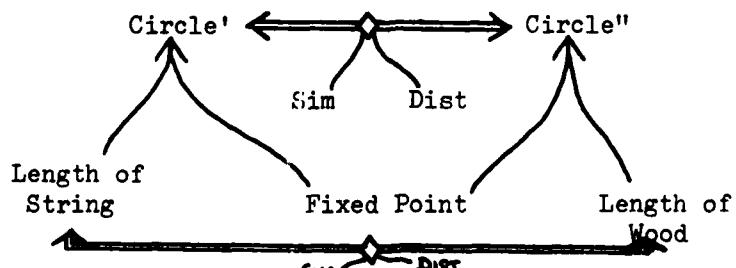


Fig 5 The resolution of Genoa

It is not correct to think that the concepts of Circle' and Circle" are analogously related. Instead the analogy relation exists between the Selprune of Circle' and the Selprune of Circle". The analogy (also referred to as a pseudonode) is, in fact, an injunction to Selprune under Circle' and Circle" at the same time ie concurrently, and to consider both Selprunes together. If this is done, then we create an analogical universe that acts as a common referent for the two prunings. A successful response to this injunction is, identically, the Lp operation, Condense. Using Figure 5, Condense is defined in the following way.

Condense (\diamond) = Superimposition (Selprune C', Selprune C")

It will be recalled from Section 2.4 that a superimposition of Selprunes gives rise to a Pruning. But what will be the head node for such a Pruning as is derived from the Condense operation (above) and where is it inscribed? To answer these questions we say that the entailment mesh produced by Adam for the conversational domain "Geometric Figures" resonates at a particular "level", Ω^0 . The Condensation of pseudonodes in mesh Ω^0 gives rise to (real) nodes in Ω' . These Ω' nodes exist in Ω^0 as Selprunes and this is true for all Ω' nodes.

Let us see what happens when Adam accepts the injunction to consider Selprune C' (in universe X) and Selprune C" (in universe Y) simultaneously (ie superimpose them) as he must do if his analogy assertion is to be useful. Essentially, he constructs a third, analogical universe, U, in which both Selprunes are represented. The analogical universe, U, is described by the manifestation of Selprunes, C' and C" in U, together with the similarities (Sim) between X & U and Y & U, and the distinctions (Dist) between X & U and Y & U. The similarities so produced give rise to the concept name in Ω' of the Condensed pseudonode of Ω^0 , which in Adam's case is Circles. Adam considers the perspectives on the domain 'Geometric Figures' that are the Selprunes Circle' and Circle" simultaneously. To do this he adopts a more general perspective (he sits in the analogical universe) and notes some distinctions (ie different Circles are produced in different ways) but he notes also from his vantage point, the qualitative similarity between Circle' and Circle" which he chooses to describe as Circles. This is the same as instating the three topics, Circles, Selprune Circle' and Selprune Circle" as a coherent bundle but in mesh Ω' , with each derivable from the other two.

The analogy in Fig 5 represents, then, an agreement (Circles) over an understanding of Selprunes, Circle' and simultaneously Circle". We have already been introduced to this idea, in section 2.2 when Adam Condensed the analogous relations between Selprunes String and Wood to form a node in Ω' he called Rigid Length of Unspecified Material. The agreement comes about as a result of a conversation that Adam has with himself spanning the two distinct universes to which the Selprunes belong. Further, this agreement is the organisational closure of the conversation. But in mesh Ω^0 the analogy only represents this agreement.

Pask (1976) refers to agreements which are represented in a mesh by analogy relations as "frozen" or "petrified" agreements, and speaks of the resuscitation of the dormant and possibly unnamed participants (eg different perspectives taken by Adam), when the analogy is understood (ie when the analogy in Ω^0 is Condensed to a node in Ω').

Though a little difficult to understand, the mechanics of analogy relations are perhaps the most important to Conversation Theory and Lp. As we have already said, the observables of Conversation theory are agreements across which information is transferred and these are represented in entailment meshes by the inscription of analogies. It is worth repeating, however, that these mechanics have been discussed here in only the broadest of terms in accordance with the stated scope of the present paper. A more detailed account of what it is for two intelligent entities to agree is given in Pask (1979).

2.6 Isomorphic Inference

The process of Isomorphic Inference is much akin to that of Saturation. In the THOUGHTSTICKER implementation of Lp, both processes are significant ways in which the eliciting system takes the initiative from the user in order to resolve questions generated by the structure of the growing representation. Indeed, the total activity of THOUGHTSTICKER is Lp.

In Fig 5, Adam halted the ambiguous oscillation of the THOUGHTSTICKER system by distinguishing two universes inhabited by Circle' and Circle" respectively. With this done, Genoa was resolved and Adam did not need to do any more. However, the process of Isomorphic Inference may ask Adam to see if he can make any further distinctions between the two universes. To do so will produce a greater stability and permanence in the analogical universe and hence give that universe more gain. (It is sometimes helpful to think of analogies as amplifiers having gain. That is, they amplify distinguishing relations between concepts that are previously embedded or only implied). One question that THOUGHTSTICKER may pose to Adam (via Isomorphic Inference) concerns the possibility of a bifurcation of Fixed Point in Fig 5. Can Adam distinguish between Fixed Point when it is used with String to derive Circle', and Fixed Point when it is used with Wood to derive Circle'? If so, then the Condensing of the newly created pseudonode will reveal previously embedded meaning in Adam's understanding of Geometric Figures from the perspective Fixed Point. If Adam had not already appreciated and instated the analogy between Selprunes, String and Wood, the process of Isomorphic Inference would suggest it.

To consider another use of the process, THOUGHTSTICKER may ask Adam as he tries to represent his understanding of, say, Ellipse, whether, by Isomorphic Inference, concepts like Fixed Point and String are appropriate. This, of course, would only occur if Adam makes the prior assertion that Ellipse and Circle are, in some ways, analogous.

3. SUMMARY

This paper is a general and incomplete examination of Pask's language Lp for the representation of conversations between intelligent entities. It is general by design in the attempt to render this esoteric system more comprehensible to readers who are new to Pask's work, yet interested to understand it.

Lp is itself an active process that represents conversations, defined as the sharing of concepts. Conversations may be held between two or more perspectives taken by one person, between two or more people or, generally, between any processors that can modulate Lp. For the concepts of one entity to be shared by another, information (not data) must be transferred and the medium of information transfer is represented in Lp by an analogy of agreement.

THOUGHTSTICKER is a computer based implementation of Lp whose activity is Lp. It is used to represent the knowledge held by one or more persons concerning any conversational domain. Knowledge is represented as topics at nodes of a mesh which may be linked by derivational, analogical or both, types of connectivity. Concepts and memories are treated by Lp as processes themselves, and are required to be coherent. Concepts are remembered, described and realised in terms of other concepts. Concepts are bundles of procedures which, on execution, result either in an image of the concept (description) or the creation of an example of the concept (behaviour). Lp therefore represents knowledge in ways which are both descriptive and prescriptive.

Derivational linking of concepts reflects coherency within and between organisationally closed bundles that represent the conversational domain. Analogical linking ensures coherence (agreement) between different perspectives on the conversational domain. Lp is complex, but then so is the problem of how intelligent entities such as ourselves ever come to understand each other. Pask has provided a major new tool for understanding this understanding, but much more than this, holds up the promise of a mechanism for non-trivial conversation between man and machines*.

*THOUGHTSTICKER is currently being implemented on AMTE/APU's computer facilities for the purpose of detailed investigation and evaluation.

REFERENCES

CLARK, P. Saturation, bifurcations and analogy in L_p : A position paper. Presented to IV Annual Richmond Conference, SRL, Richmond (1979)

ENTWHISTLE, N.J. Symposium: knowledge structures and styles of learning: A summary of Pask's recent research. Brit J. Ed Psychol, 48 (1978)

PASK, G. Conversation, cognition and learning. Amsterdam & New York, Elsevier (1975)

PASK, G. Conversation theory: applications in education and epistemology. Amsterdam & New York, Elsevier (1976)

PASK, G. A protolanguage: (L_p , the THOUGHTSTICKER language) SRL, Richmond (1979)

PASK, G. (in press) (Developments in conversation theory) to be published in Int J Man Machine Studies

ACKNOWLEDGEMENTS

The author wishes to acknowledge the contribution made by Colin Sheppard of AMTE/APU to the preparation of the seminars and for his comments on the drafts of these papers.

Particular thanks also go to Paul Pangaro of SRL for his part in our most stimulating discussion of L_p .

GENERAL REFERENCES

BOXER, P.J. Reflective analysis. *Int J Man Machine Studies (IJMMS)*, 11, 547-584 (1979)

BURTON, R. R. & BROWN, J.S. An investigation of computer coaching for informal learning activities. *IJMMS*, 11, 5-24 (1979)

CARBONELL, J.G. Subjective understanding: computer models of belief systems. *Yale Uni. Dept Comp Sciences. Res Rep 150* (1979)

CARR, B. Wumpus advisor II. *MIT AI Lab. AIM 417 Massachusetts Institute of Technology, Cambridge, Mass* (1977)

CARR, B & GOLDSTEIN, I.P. Overlays: A theory of modelling for computer aided instruction. *MIT AI Lab. AIM 406. AD-A052 210* (1977)

CLANCY, W.J. Tutoring rules for guiding a case method dialogue *IJMMS*, 11, 25-49 (1979)

ENTWHISTLE, N.J. Symposium: knowledge structures and styles of learning: A summary of Pask's recent research. *Br J Educ Psychol*, 48, 1-11 (1978)

GAINES, B. Foundations of fuzzy reasoning, *IJMMS*, 8, 623-668 (1976)

GAINES, B. Man computer communication - What next? *IJMMS*, 10, 225-232 (1978)

GOLDSTEIN, I.P. The genetic graph: A representation for the evolution of procedural knowledge. *IJMMS*, 11, 51-77 (1979)

GOLDSTEIN, I.P. & CARR, B. The computer as coach: An athletic paradigm for intellectual education. *Proceedings of 1977 Annual Conference Association for Computing Machinery, Seattle* (1977)

HAYES, P. Trends in artificial intelligence, *IJMMS*, 10, 295-299 (1978)

KNOOP, P.A. Man computer symbiosis through interactive graphics: A survey and identification of critical research areas. *Advanced System Division, AFHRL, Ohio* (1977)

LESLING, P.D., BLANK, M.S. & ANTERSTON, T.A. ZORK: A computerised fantasy game, *MIT Lab Comp Science* (1979)

MILLER, M.L. A structured planning and debugging environment for elementary programming, *IJMMS*, 11, 79-95 (1978)

O'SHEA, T. A self-improving quadratic tutor, *IJMMS*, 11, 97-124 (1978)

PASK, G. A protolanguage (L : the THOUGHTSTICKER language) System Research Ltd, Woodville House, Richmond, Surrey (1979)

PASK, G. Users guide for the THOUGHTSTICKER system. SRL Richmond Surrey (1979)

PASK, G. An essay on the kinetics of language, behaviour and thought. Proceedings Silver Anniversary Int Meeting Society for Gen Systems Research, London (1979)

PASK, G. Automation, coherence in organisations and people who form a part of them. Infotech International (1979)

PASK, G. The limits of togetherness. SRL Richmond Surrey (1980)

PASK, G. (in press)(Developments in conversation theory) to be published in IJMMS.

SELF, J.A. Student models in computer aided instruction. IJMMS, 6 261-276 (1974)

SHAW, M.L.G. Conversational heuristics for eliciting shared understanding, IJMMS 11, 621-634 (1979)

SLEEMAN, D.G. & HENDLEY, R.J. ACE: A system which analyses complex explanations, IJMMS 11, 125-144 (1979)

STEVENS, A.L. & COLLINS, A. The goal structure of a socratic tutor. Proceedings of 1977 Annual Conference, Assoc Computing Machinery, Seattle (1977)

STEVENS, A.L. COLLINS, A. & GOLDIN, S.E. Misconception in student's understanding, IJMMS, 11, 145-156 (1979)

WATERMAN, D.A. A rule based approach to knowledge acquisition for man-machine interface programs, IJMMS, 10 693-711 (1978)

WALTZ, D.L. An English language question answering system for a large relational database. Comms of the ACM Vol 21, No 7 (1978)

This list of references is by no means exhaustive. Since February 1980 our attention has been drawn to many other publications. The present list represents the starting point for APU's Man Computer Studies Group.

